



CHAPTER 7

Circuits and Tunnels

This chapter explains Cisco ONS 15600 SDH high-order circuits, data communications channel (DCC), and IP-encapsulated tunnels. To provision circuits and tunnels, refer to the *Cisco ONS 15600 SDH Procedure Guide*.

Chapter topics include:

- [7.1 Overview, page 7-2](#)
- [7.2 Circuit Properties, page 7-2](#)
- [7.3 Cross-Connect Card Bandwidth, page 7-11](#)
- [7.4 DCC Tunnels, page 7-11](#)
- [7.5 Multiple Destinations for Unidirectional Circuits, page 7-12](#)
- [7.6 SNCP Circuits, page 7-12](#)
- [7.7 Protection Channel Access Circuits, page 7-14](#)
- [7.8 MS-SPRing VC4/VC3 Squelch Table, page 7-15](#)
- [7.9 Section and Path Trace, page 7-15](#)
- [7.10 Automatic Circuit Routing, page 7-16](#)
- [7.11 Manual Circuit Routing, page 7-17](#)
- [7.12 Constraint-Based Circuit Routing, page 7-19](#)
- [7.13 Bridge and Roll, page 7-19](#)
- [7.14 Merged Circuits, page 7-24](#)
- [7.15 Reconfigured Circuits, page 7-25](#)
- [7.16 Server Trails, page 7-25](#)
- [7.17 Traffic Routing over a Third-Party Network, page 7-26](#)
- [7.18 Low-Order Traffic Routing over an ONS 15600 SDH Hub Node, page 7-27](#)
- [7.19 High-Order VC3 Traffic Routing, page 7-29](#)



Note

In this chapter, “cross-connect” and “circuit” have the following meanings: cross-connect refers to the connections that occur within a single ONS 15600 SDH to allow a circuit to enter and exit an ONS 15600 SDH. Circuit refers to the series of connections from a traffic source (where traffic enters the ONS 15600 SDH network) to the destination (where traffic exits an ONS 15600 SDH network).

7.1 Overview

You can create circuits across and within ONS 15600 SDH 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.
- 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 15600 SDH subnetwork connection protection (SNCP) ring with third-party equipment SNCPs.
- Set SNCP circuits as revertive or nonrevertive.

You can provision circuits at any of the following points:

- Before cards are installed. The ONS 15600 SDH allows you to provision slot 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 enabled. Circuits do not actually carry traffic until the cards and SFPs are installed and the ports are in the Unlocked-enabled state; the Locked-enabled, maintenance state; or the Unlocked-disabled, automatic InService state. Circuits carry traffic as soon as the signal is received.

7.2 Circuit Properties

The ONS 15600 SDH Circuits window ([Figure 7-1 on page 7-3](#)), which is available from network, node, and card view, is where you can view information about circuits, including:

- Name—The name of the circuit (user-assigned or automatically generated).
- Type—For the ONS 15600 SDH, the circuit type is HOP (high-order circuit).
- Size—The circuit size can be VC3, VC4, VC4-4c, VC4-8c, VC4-16c, or VC4-64c. ASAP optical ports also allow circuit sizes of VC4-2c and VC4-3c. For time slot availability on concatenated VCs, see the [“7.2.1 Concatenated VC4 Time Slot Assignments” section on page 7-4](#).
- Protection—The protection type; see the [“7.2.4 Circuit Protection Types” section on page 7-8](#).
- Direction—The circuit direction, either two-way or one-way.
- Status—The circuit status; for details, see the [“7.2.2 Circuit Status” section on page 7-6](#).
- Source—The circuit source in the format: *node/slot/port/virtual container*. If an ASAP PPM port is the circuit source, the port format is *PIM-PPM-port*, where pluggable interface module (PIM) and PPM values are 1 through 4 (for example, p1-1-1). PPMs have only one port.
- Destination—The circuit destination in the format: *node/slot/port/virtual container*. If an ASAP PPM port is the circuit destination, the port format is *PIM-PPM-port*, where PIM and PPM values are 1 through 4 (for example, p1-1-1). PPMs have only one port.

- # of VLANs—(Future use) 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/virtual container*.
- State—The circuit state; for details, see the “7.2.3 Circuit States” section on page 7-7.

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 7-1 ONS 15600 SDH Circuit Window in Network View

Circuit Name	Type	Size	OCHNC Wlen	Dir	Protection	Status	Source
Unknown	HOP	VC4	N/A	2-way	N/A	DISCOVERED_TL1	CXC66/s3/p1/vc4-2
Unknown	HOP	VC4	N/A	2-way	N/A	DISCOVERED_TL1	CXC66/s3/p1/vc4-1

The Circuit Filter feature allows you to reduce the number of circuits that appear in the Circuits window. You can specify certain filter criteria, such as name, direction, and state; only the circuits that match the criteria will appear in the Circuits window.



Note

You cannot set up low-order circuits to terminate on an ONS 15600 SDH node. However, you can create both high-order and low-order circuits that have an ONS 15454 SDH source and destination with an ONS 15600 SDH as a pass-through node. For information, see the “7.18 Low-Order Traffic Routing over an ONS 15600 SDH Hub Node” section on page 7-27.

7.2.1 Concatenated VC4 Time Slot Assignments

Table 7-1 shows the available time slot assignments for concatenated VC4s when using Cisco Transport Controller (CTC) to provision circuits.

Table 7-1 VC4 Mapping Using CTC

Starting VC4	VC4	VC4-2c	VC4-3c	VC4-4c	VC4-6c	VC4-8c	VC4-12c	VC4-16c	VC4-64c
1	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
3	Yes	Yes	No	No	Yes	Yes	Yes	No	No
4	Yes	No	Yes	No	Yes	Yes	Yes	No	No
5	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
6	Yes	Yes	Yes	No	Yes	Yes	No	No	No
7	Yes	Yes	Yes	No	Yes	Yes	No	No	No
8	Yes	No	No	No	Yes	Yes	No	No	No
9	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
10	Yes	Yes	Yes	No	Yes	No	No	No	No
11	Yes	Yes	No	No	Yes	No	No	No	No
12	Yes	No	No	No	No	No	No	No	No
13	Yes	Yes	Yes	Yes	Yes	No	Yes	No	No
14	Yes	Yes	Yes	No	No	No	No	No	No
15	Yes	Yes	No	No	No	No	No	No	No
16	Yes	No	Yes	No	No	No	No	No	No
17	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
18	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
19	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
20	Yes	No	No	No	Yes	Yes	Yes	No	No
21	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
22	Yes	Yes	Yes	No	Yes	Yes	No	No	No
23	Yes	Yes	No	No	Yes	Yes	No	No	No
24	Yes	No	No	No	Yes	Yes	No	No	No
25	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
26	Yes	Yes	Yes	No	Yes	No	No	No	No
27	Yes	Yes	No	No	Yes	No	No	No	No
28	Yes	No	Yes	No	No	No	No	No	No
29	Yes	Yes	Yes	Yes	No	No	No	No	No
30	Yes	Yes	Yes	No	No	No	No	No	No
31	Yes	Yes	Yes	No	Yes	No	No	No	No

Table 7-1 VC4 Mapping Using CTC (continued)

Starting VC4	VC4	VC4-2c	VC4-3c	VC4-4c	VC4-6c	VC4-8c	VC4-12c	VC4-16c	VC4-64c
32	Yes	No	No	No	No	No	No	No	No
33	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
34	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
35	Yes	Yes	No	No	Yes	Yes	Yes	No	No
36	Yes	No	No	No	Yes	Yes	Yes	No	No
37	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
38	Yes	Yes	Yes	No	Yes	Yes	No	No	No
39	Yes	Yes	No	No	Yes	Yes	No	No	No
40	Yes	No	Yes	No	Yes	Yes	No	No	No
41	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
42	Yes	Yes	Yes	No	Yes	No	No	No	No
43	Yes	Yes	Yes	No	Yes	No	No	No	No
44	Yes	No	No	No	No	No	No	No	No
45	Yes	Yes	Yes	Yes	No	No	No	No	No
46	Yes	Yes	Yes	No	No	No	No	No	No
47	Yes	Yes	No	No	No	No	No	No	No
48	Yes	No	No	No	No	No	No	No	No
49	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
50	Yes	Yes	Yes	No	Yes	Yes	Yes	No	No
51	Yes	Yes	No	No	Yes	Yes	Yes	No	No
52	Yes	No	Yes	No	Yes	Yes	Yes	No	No
53	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No
54	Yes	Yes	Yes	No	Yes	Yes	No	No	No
55	Yes	Yes	Yes	No	Yes	Yes	No	No	No
56	Yes	No	No	No	Yes	Yes	No	No	No
57	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
58	Yes	Yes	Yes	No	Yes	No	No	No	No
59	Yes	Yes	No	No	Yes	No	No	No	No
60	Yes	No	No	No	No	No	No	No	No
61	Yes	Yes	Yes	Yes	Yes	No	No	No	No
62	Yes	Yes	Yes	No	Yes	No	No	No	No
63	Yes	Yes	No	No	Yes	No	No	No	No
64	Yes	No	No	No	Yes	No	No	No	No

7.2.2 Circuit Status

The circuit statuses that appear in the Circuit window Status column are generated by CTC based on an assessment of conditions along the circuit path. [Table 7-2](#) shows the statuses that can appear in the Status column.

Table 7-2 ONS 15600 SDH Circuit Status

Status	Definition/Activity
CREATING	CTC is creating a circuit.
DISCOVERED	CTC created a circuit. All components are in place and a complete path exists from circuit source to destination.
DELETING	CTC is deleting a circuit.
PARTIAL	<p>A CTC-created circuit is missing a connection or circuit span (network link), 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 Transaction Language One (TL1-created) circuit or a TL1-like, CTC-created circuit is complete. A complete path from source to destination(s) exists.
PARTIAL_TL1	A TL1-created circuit or a TL1-like, CTC-created circuit is missing a cross-connect or circuit span (network link), and a complete path from source to destination(s) does not exist.
CONVERSION_PENDING	An existing circuit in a topology upgrade is set to this state. The circuit returns to the DISCOVERED state when the topology upgrade is complete. For more information about topology upgrades, see Chapter 8, “SDH Topologies and Upgrades.”

Table 7-2 ONS 15600 SDH 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 it is a temporary circuit. These circuits can be deleted if a topology upgrade fails. For more information about topology upgrades, see Chapter 8, “SDH Topologies and Upgrades.”
DROP_PENDING	A circuit is set to this status when a new circuit drop is being added.

7.2.3 Circuit States

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

- If all cross-connects in a circuit are in the Unlocked-enabled service state, the circuit service state is Unlocked.
- If all cross-connects in a circuit are in the Locked-enabled,maintenance, Locked-enabled,disabled service state, or the Unlocked-disabled,automaticInService state, the circuit service state is Locked.
- Partial is appended to the Locked circuit service state when circuit cross-connect state are mixed and not all in the Unlocked-enabled service state. The Locked-partial state can occur during automatic or manual transitions between states. The Locked-partial service state can appear during a manual transition caused by an abnormal event such as a CTC crash or communication error, or if one of the cross-connects could not be changed. Refer to the *Cisco ONS 15600 SDH Troubleshooting Guide* for troubleshooting procedures. The Locked-partial circuit state does not apply to optical channel network connection (OCHNC) circuit types.

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

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



Note

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

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

Circuits do not use the soak timer, but ports do. The soak period is the amount of time that the port remains in the Unlocked-disabled,automaticInService service state after a signal is continuously received. When the cross-connects in a circuit are in the Unlocked-disabled,automaticInService service state, the ONS 15600 SDH monitors the cross-connects for an error-free signal. It changes the state of the circuit from Locked to Unlocked or to Locked-partial as each cross-connect assigned to the circuit path is completed. This allows you to provision a circuit using TL1, verify its path continuity, and prepare the port to go into service when it receives an error-free signal for the time specified in the port soak timer.

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

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

7.2.4 Circuit Protection Types

The Protection column in the Circuit window shows the card (line) and SDH topology (path) protection used for the entire circuit path. [Table 7-3](#) lists the protection type indicators that you will see in this column.

Table 7-3 *Circuit Protection Types*

Protection Type	Description
1+1	The circuit is protected by a 1+1 protection group.
2F MS-SPRing	The circuit is protected by a two-fiber MS-SPRing.
2F-PCA	The circuit is routed on a protection channel access (PCA) path on a two-fiber MS-SPRing. PCA circuits are unprotected.
DRI	The circuit is protected by dual-ring interconnect (DRI).
N/A	A circuit with connections on the same node is not protected.
Protected	The circuit is protected by diverse SDH topologies, for example, an MS-SPRing and an SNCP, or an SNCP 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.
Unprot (red)	A circuit created as a fully protected circuit is no longer protected due to a system change, such as removal of an MS-SPRing or 1+1 protection group.
SNCP	The circuit is protected by an SNCP.

7.2.5 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 [“7.5 Multiple Destinations for Unidirectional Circuits”](#) section on page 7-12.
- **SNCP Selectors**—Allows you to change SNCP selectors. For more information, see the [“7.6 SNCP Circuits”](#) section on page 7-12.
- **SNCP Switch Counts**—Allows you to change SNCP switch protection paths. For more information, see the [“7.6 SNCP Circuits”](#) section on page 7-12.

- State—Allows you to edit cross-connect service states.
- Merge—Allows you to merge aligned circuits. For more information, see the “7.14 Merged Circuits” section on page 7-24.

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

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

- The circuit direction (unidirectional or bidirectional)
- The node names
- The slot and port numbers
- The circuit source and drop

The pane area shows:

- The circuit name, type, and size
- The protection type (SNCP, unprotected, MS-SPRing, 1+1) and routing preferences such as fully protected path
- Open Shortest Path First (OSPF) area IDs

For MS-SPRings, the detailed map shows the number of MS-SPRing fibers and the MS-SPRing ring ID. For SNCPs, the map shows the active and standby paths from circuit source to destination, and it also shows the working and protect paths. Selectors appear as pentagons on the detailed circuit map.

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 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, respectively. Port states are indicated by colors, shown in [Table 7-4](#).

Table 7-4 Port State Color Indicators

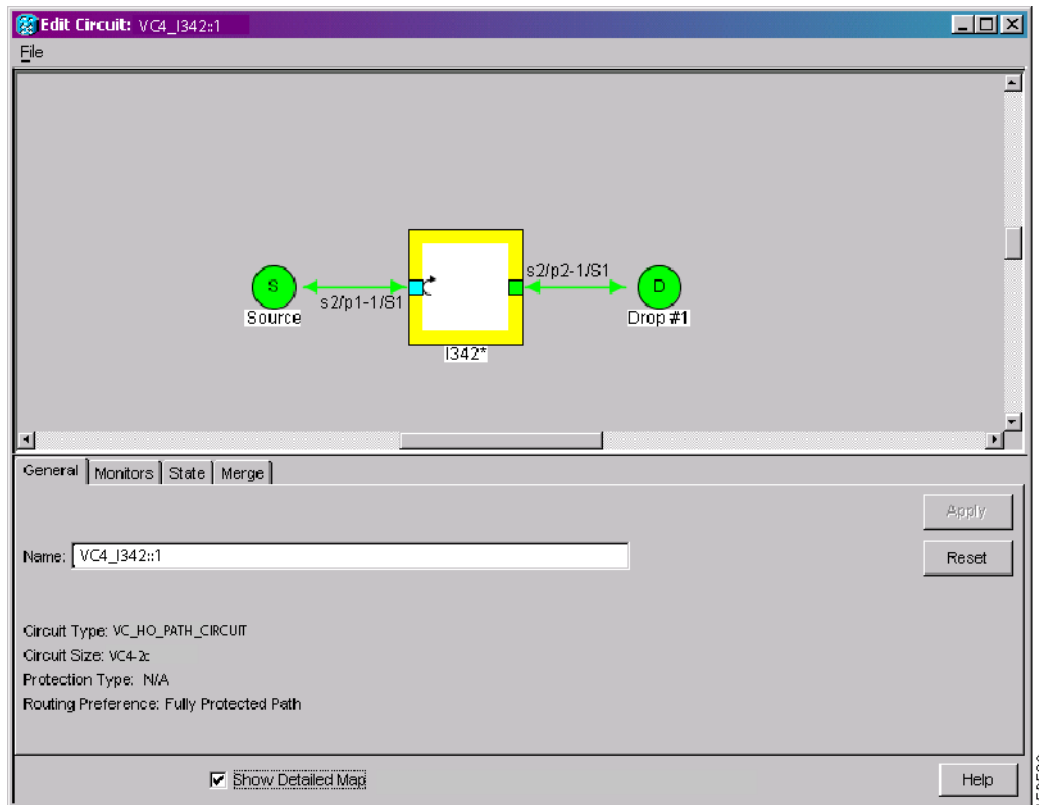
Port Color	Service State
Green	Unlocked-enabled
Gray	Locked-enabled,disabled
Violet	Unlocked-disabled,automaticInService
Blue (Cyan)	Locked-enabled,maintenance

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

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

Figure 7-2 shows an example of the Edit Circuit window with a terminal loopback.

Figure 7-2 Detailed Circuit Map Showing a Terminal 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), 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 an SNCP span to change the state of the path selectors in the SNCP circuit.

7.3 Cross-Connect Card Bandwidth

The single shelf cross-connect card (SSXC) performs port-to-port time-division multiplexing (TDM). The VC matrix has the capacity for 2048 VC4 and/or 6144 VC3 terminations. Because each VC4/VC3 circuit requires a minimum of two terminations, one for ingress and one for egress, the SSXC has a capacity for 1024 VC4 and 3072 VC3 circuits respectively. However, this capacity is reduced at SNCP and 1+1 nodes because three VC4/VC3 terminations are required at circuit source and destination nodes and four terminations are required at 1+1 circuit pass-through nodes. SNCP pass-through nodes only require two VC4/VC3 terminations.

7.4 DCC Tunnels

SDH provides four data communications channels (DCCs) for network element operations, administration, maintenance, and provisioning (OAM&P): one on the SDH Section layer (DCC1) and three on the SDH Line layer (DCC2, DCC3, DCC4). The ONS 15600 SDH uses the regenerator-section DCC (RS-DCC) for ONS 15600 SDH management and provisioning. An RS-DCC and multiplex-section (MS-DCC) each provide 192 Kbps of bandwidth per channel. The aggregate bandwidth of the three MS-DCCs is 576 Kbps. You can tunnel third-party SDH equipment across ONS 15600 SDH networks using one of two tunneling methods, a traditional DCC tunnel or an IP-encapsulated tunnel.

7.4.1 Traditional DCC Tunnels

In traditional DCC tunnels, you can use the three MS-DCCs and the RS-DCC (when not used for ONS 15600 SDH DCC terminations). A DCC tunnel endpoint is defined by slot, port, and DCC, where DCC can be either the RS-DCC or one of the MS-DCCs. You can link MS-DCCs to MS-DCCs and link RS-DCCs to RS-DCCs. You can also link an RS-DCC to an MS-DCC and an MS-DCC to an RS-DCC. To create a DCC tunnel, connect the tunnel endpoints from one ONS 15600 SDH optical port to another. [Table 7-5](#) lists the DCC tunnels that you can create.

Table 7-5 DCC Tunnels

DCC	SDH Layer	SDH Bytes
DCC1	RS-DCC	D1 to D3
DCC2	MS-DCC	D4 to D6
DCC3	MS-DCC	D7 to D9
DCC4	MS-DCC	D10 to D12

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

- Each ONS 15600 SDH can have up to 64 DCC tunnel connections.
- An RS-DCC that is terminated cannot be used as a DCC tunnel endpoint.
- An RS-DCC that is used as a DCC tunnel endpoint cannot be terminated.
- All DCC tunnel connections are bidirectional.

7.4.2 IP-Encapsulated Tunnels

An IP-encapsulated tunnel puts an RS-DCC in an IP packet at a source node and dynamically routes the packet to a destination node. 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 RS-DCC bandwidth.

Each ONS 15600 SDH supports up to 128 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.

7.5 Multiple Destinations for Unidirectional Circuits

Unidirectional circuits can have multiple destinations (drops) for use in broadcast circuit schemes. In broadcast scenarios, one source transmits traffic to multiple destinations, but traffic is not returned to the source. The ONS 15600 SDH supports one of the following:

- Up to 2048 1:2 nonblocking broadcast connections
- Up to 682 1: n nonblocking broadcast connections (where n is less than or equal to 8)

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. To create an alarm profile, refer to the “Manage Alarms” chapter in the *Cisco ONS 15600 SDH Procedure Guide*.

7.6 SNCP Circuits

Use the Edit Circuit window to change SNCP selectors and switch protection paths (Figure 7-3). In the SNCP Selectors subtab in the Edit Circuits window, you can:

- View the SNCP circuit’s working and protection paths.
- Edit the reversion time.
- Set the hold-off timer.
- Edit the Signal Fail (SF) and Signal Degrade (SD) thresholds.



Note

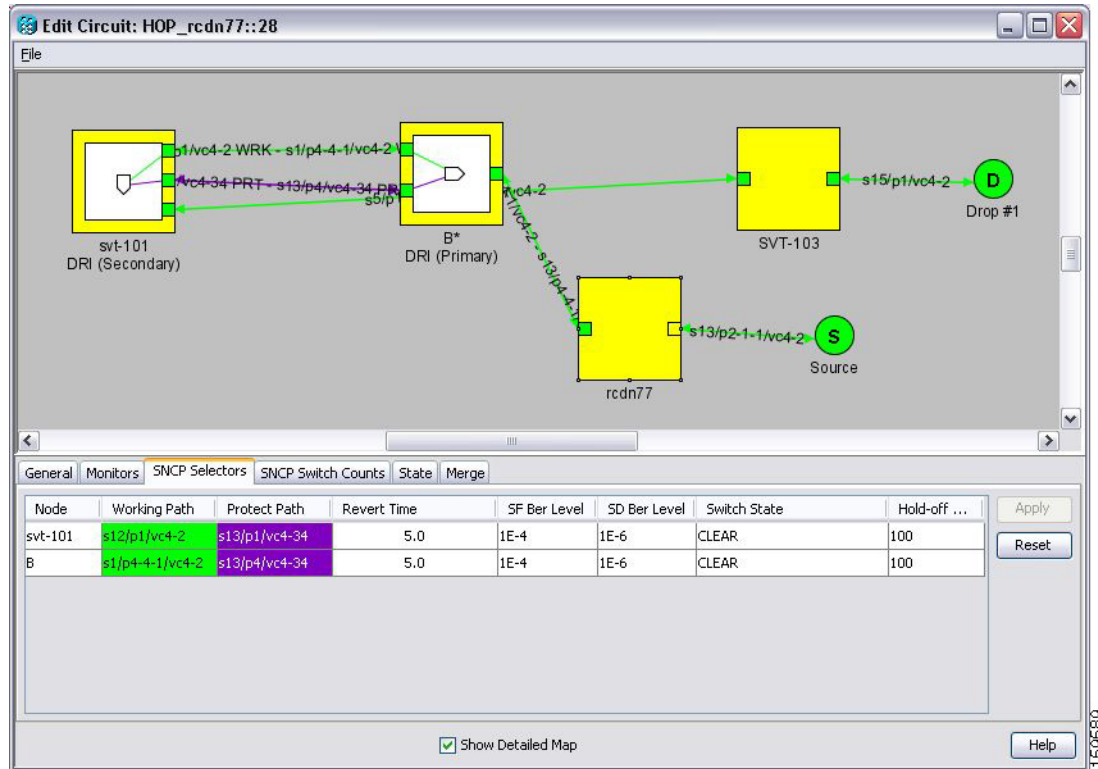
In the SNCP Selectors tab, the SF Ber Level and SD Ber Level columns display “N/A” for those nodes that do not support low-order path signal bit error rate (BER) monitoring.

From the SNCP Switch Counts subtab, you can:

- Perform maintenance switches on the circuit selector.

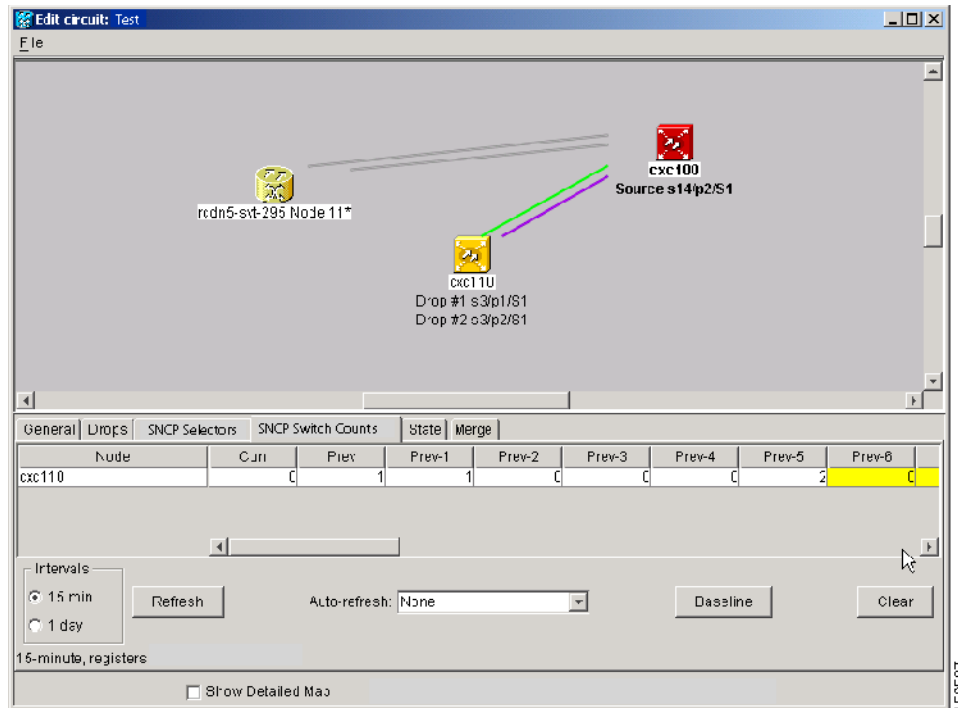
- View switch counts for the selectors.

Figure 7-3 Editing SNCP Selectors



On the SNCP Switch Counts tab, you can view switch counts for the selectors (Figure 7-4).

Figure 7-4 Viewing SNCP Switch Counts



In network view, you can right-click on a span and choose the Circuits command to open the Circuits on Span dialog box. Use the Circuits on Span dialog box to view circuit span usage information and, for SNCP spans, to perform span switches and lockouts.

7.7 Protection Channel Access Circuits

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

Traffic provisioning on MS-SPRing 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 15600 SDH Procedure Guide* for more information. When provisioning PCA circuits, two considerations are important to keep in mind:

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

7.8 MS-SPRing VC4/VC3 Squelch Table

The ONS 15600 SDH platform does not support low-order path squelching; however, when an ONS 15454 SDH and an ONS 15600 SDH are in the same network, the ONS 15600 SDH node allows the ONS 15454 SDH node to carry low-order path circuits in a VC_LO_PATH_TUNNEL. The ONS 15600 SDH performs 100-ms VC4-level squelching for each low-order-access VC at the switching node in case of a node failure.

- SDH-AU4—Displays all circuits with SDH-AU4 format (VC4 or VC4-*nc* high-order circuits). For example, if “Circuit1” is a high-order VC4 circuit using ring bandwidth VC4-2c, then “Circuit1” appears in the SDH-AU4 pane in the column “vc4=2.”
- SDH-AU3—Displays all circuits with SDH-AU3 format (VC3 high-order circuits).
- VC4/VC3 Number—The SDH-AU4 pane shows VC4 number only, and the SDH-AU3 pane shows VC3 number only.
- West Source—If traffic is received by the node on its west span, the MS-SPRing node ID of the source appears. (To view the MS-SPRing node IDs for all nodes in the ring, click the Ring Map button.)
- West Dest—If traffic is sent on the node’s west span, the MS-SPRing node ID of the destination appears.
- East Source—If traffic is received by the node on its east span, the MS-SPRing node ID of the source appears.
- East Dest—If traffic is sent on the node’s east span, the MS-SPRing node ID of the destination appears.
- Retrieve—Retrieves the most current squelch table display.



Note

MS-SPRing squelching is performed on VC4s that carry VC4 circuits only or on VC3s that carry VC3 circuits only.

7.9 Section and Path Trace

SDH J0 section and J1 path trace are repeated, fixed-length strings used to monitor interruptions or changes to circuit traffic. J0 section trace includes 1 or 16 consecutive J0 bytes. All ONS 15600 SDH optical cards and ports support J0 path trace. J1 path trace includes 16 consecutive J1 bytes. If the string received at a circuit drop port does not match the string the port expects to receive, the Trace Identifier Mismatch Path (TIM-P) alarm is raised.

Table 7-6 lists the ONS 15600 SDH cards that support J1 path trace.

Table 7-6 ONS 15600 SDH Cards Supporting J1 Path Trace

Card	Receive	Transmit
OC48/STM16 SR/SH 16 Port 1310	Yes	No
OC48/STM16 LR/LH 16 Port 1550	Yes	No
OC192/STM64 SR/SH 4 Port 1310	Yes	No
OC192/STM64 LR/LH 4 Port 1550	Yes	No
OC192LR/STM64 4 Port ITU C-Band	Yes	No

Table 7-6 ONS 15600 SDH Cards Supporting J1 Path Trace (continued)

Card	Receive	Transmit
ASAP STM-N ports	Yes	No
ASAP Ethernet ports	Yes	Yes

The ONS 15600 SDH supports both automatic and manual J1 path trace monitoring to detect and report the contents of the 16-byte path trace message (nonterminated) for the designated path.

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

**Note**

When J1 path trace is enabled on a two-fiber MS-SPRing circuit, CTC will not retrieve the path trace information from the card view Maintenance > Path Trace tab.

7.10 Automatic Circuit Routing

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

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.
- 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.

7.10.1 Bandwidth Allocation and Routing

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


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

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

7.10.2 Secondary Sources and Destination

CTC supports secondary sources and destinations (drops). Secondary sources and destinations typically interconnect two third-party networks, as shown in [Figure 7-5](#). Traffic is protected while it traverses a network of ONS 15600 SDHs.

Figure 7-5 Secondary Sources and Drops



Several rules apply to secondary sources and destinations:

- CTC does not allow a secondary destination for unidirectional circuits because you can specify additional destinations after you create the circuit.
- Primary and secondary sources should be on the same node.
- Primary and secondary destinations should be on the same node.
- Secondary sources and destinations are permitted only for regular VC4 connections.

For bidirectional circuits, CTC creates an SNCP connection at the source node that allows traffic to be selected from one of the two sources on the ONS 15600 SDH network. If you check the Fully Path Protected option during circuit creation, traffic is protected within the ONS 15600 SDH network. At the destination, another SNCP connection is created to bridge traffic from the ONS 15600 SDH network to the two destinations. A similar but opposite path exists for the reverse traffic flowing from the destinations to the sources.

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



Note

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

7.11 Manual Circuit Routing

Routing circuits manually allows you to:

- Choose a specific path, not just the shortest path chosen by automatic routing.
- Choose a specific VC4 on each link along the route.

CTC imposes the following rules on manual routes:

- All circuits in a shared packet ring should have links with a direction that flows from source to destination.
- If you enabled Fully Protected Path, choose a diverse protect (alternate) path for every unprotected segment.
- For a node that has an SNCP selector based on the links chosen, the input links to the SNCP selectors cannot be 1+1 protected. The same rule applies at the SNCP bridge.

If you enabled Fully Protected Path, 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 7-7 summarizes the available bidirectional connections. Any other combination is invalid and generates an error.

Table 7-7 Bidirectional VC4 Circuits

No. of Inbound Links	No. of Outbound Links	No. of Sources	No. of Drops	Connection Type
—	2	1	—	SNCP
2	—	—	1	SNCP
2	1	—	—	SNCP
1	2	—	—	SNCP
1	—	—	2	SNCP
—	1	2	—	SNCP
2	2	—	—	Double SNCP
2	—	—	2	Double SNCP
—	2	2	—	Double SNCP
1	1	—	—	Two-way

Table 7-8 summarizes the available unidirectional connections. Any other combination is invalid and generates an error.

Table 7-8 Unidirectional VC4 Circuits

No. of Inbound Links	No. of Outbound Links	No. of Sources	No. of Drops	Connection Type
1	1	—	—	One-way
1	2	—	—	SNCP head end
—	2	1	—	SNCP head end
2	—	—	1+	SNCP drop and continue

7.12 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 comprised entirely of MS-SPRing or 1+1 links, or the path between source and destination can be entirely protected using 1+1 or MS-SPRing links, no extended SNCP, or virtual SNCP, protection is used.

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

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

When you choose automatic circuit routing during circuit creation, you have the option to require or exclude nodes and links in the calculated route. You can use this option to 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 are either free or use less bandwidth. By selecting a required node and/or a link, you force CTC to use (or not use) an element, resulting in more efficient use of network resources.

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

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

7.13 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 SDH, ONS 15600, ONS 15454, ONS 15454 SDH, ONS 15310-MA, and ONS 15310-CL.

7.13.1 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 7-6](#) shows the Rolls window.

Figure 7-6 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 “7.13.2 Roll Status” section on page 7-21 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 once 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.

7.13.2 Roll Status

Table 7-9 lists the roll statuses. You can only reroute circuits that have a DISCOVERED status. (See Table 7-2 on page 7-6 for a list of circuit statuses.) You cannot reroute circuits that are in the ROLL_PENDING status.

Table 7-9 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.

7.13.3 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 7-7 on page 7-22 and Figure 7-8 on page 7-22, respectively).
- Roll a segment of the circuit onto another chosen circuit (Figure 7-9 on page 7-22). This roll also results in a new destination or a new source.


In Figure 7-7, you can select any available VC4 on Node 1 for a new source.

Figure 7-7 *Single Source Roll*


In [Figure 7-8](#), you can select any available VC4 on Node 2 for a new destination.

Figure 7-8 *Single Destination Roll*

[Figure 7-9](#) 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 7-9 *Single Roll from One Circuit to Another Circuit (Destination Changes)*

[Figure 7-10](#) shows one circuit rolling onto another circuit at the source.

Figure 7-10 *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 7-11 illustrates a dual roll on the same circuit.

Figure 7-11 Dual Roll to Reroute a Link





Figure 7-12 illustrates a dual roll involving two circuits.

Figure 7-12 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.

7.13.4 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.

7.13.5 Protected Circuits

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

7.14 Merged Circuits

A circuit merge combines a single selected circuit with one or more circuits. You can merge CTC-created circuits and TL1-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 Circuit 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.
- 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 end points must send or receive the same framing format.
- The merged circuits must become a DISCOVERED circuit.

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

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

7.15 Reconfigured Circuits

You can reconfigure multiple circuits, which is typically necessary when a large number of circuits are in the PARTIAL status. When you reconfigure multiple circuits, the selected circuits can be any combination of DISCOVERED, PARTIAL, DISCOVERED_TL1, and PARTIAL_TL1 circuits. You can reconfigure tunnels, 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.

7.16 Server Trails

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

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

**Note**

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

7.16.1 Server Trail Protection Types

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

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

**Note**

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

7.16.2 VCAT Circuit Routing over Server Trails

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

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



Note

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

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

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



Note

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

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

7.16.2.1 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.

7.17 Traffic Routing over a Third-Party Network

If ONS 15600 SDHs are connected to a third-party network, you can create an open-ended SNCP circuit to route a circuit through it. To do this, you create three circuits. One circuit is created on the source ONS 15600 SDH network. This circuit has one source and two destinations, one at each ONS 15600 SDH that is connected to the third-party network. The second circuit is created on the third-party network so that the circuit travels across the network on two paths to the ONS 15600 SDHs. That circuit routes the two circuit signals across the network to ONS 15600 SDHs that are connected to the network on other side. At the destination node network, the third circuit is created with two sources, one at each node connected to the third-party network. A selector at the destination node chooses between the two signals that arrive at the node, similar to a regular SNCP circuit.

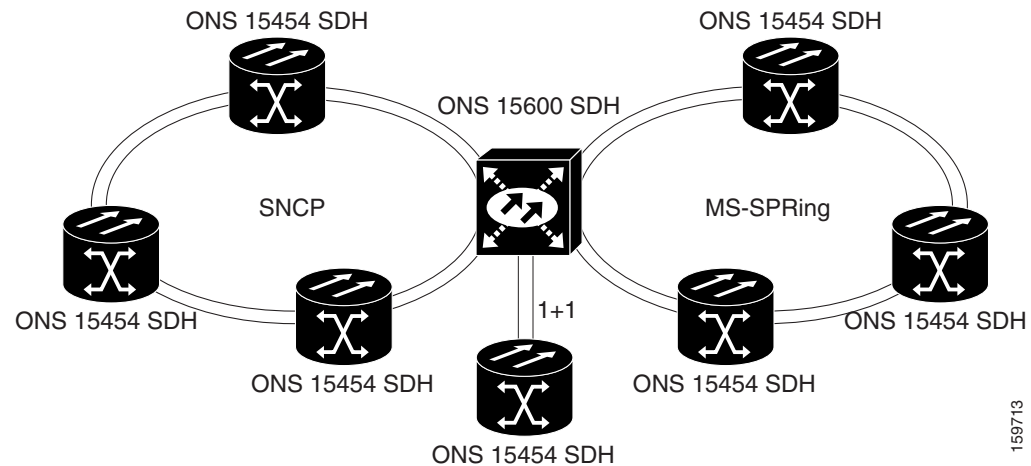
7.18 Low-Order Traffic Routing over an ONS 15600 SDH Hub Node

When an ONS 15600 SDH node is configured as a hub between compatible protection domains, it is possible to use CTC to automatically create circuits, including low-order circuits from an ONS 15454 SDH node, across the domains. Low-order circuits cannot terminate on an ONS 15600 SDH node. All combinations of protection domains are compatible for high-order circuits, including combinations of SNCP and line-protected domains.

When you must route low-order traffic over an ONS 15600 SDH hub node connecting SNCP, line-protected, and unprotected domains, you must manually create the circuits for this traffic. For specific information on ONS 15454 SDH low-order circuit creation and tunneling, refer to the circuit chapters in the *Cisco ONS 15454 SDH Reference Manual*.

Figure 7-13 shows an ONS 15600 SDH as a protection domain hub between a ONS 15454 SDH SNCP network, an ONS 15454 SDH MS-SPRing network, and a 1+1 protected ONS 15454 SDH node.

Figure 7-13 ONS 15600 SDH as Hub Node between Protection Domains



7.18.1 Automatic Low-Order Circuit Creation

To route low-order traffic across an ONS 15600 SDH for compatible protection domains, you must create a low-order tunnels on an ONS 15454 SDH node or create an end-to-end circuit. If you create an end-to-end circuit, CTC automatically creates a tunnel to route the low-order traffic. You can create circuits or tunnels with an ONS 15600 SDH as a hub node for the following domain combinations:

- SNCP to SNCP
- 1+1 to 1+1
- MS-SPRing to MS-SPRing
- 1+1 to MS-SPRing
- Unprotected to MS-SPRing

Once the tunnel or end-to-end circuit exists, additional low-order circuits can be routed until all of the bandwidth is used. The tunnel source and destination are the low-order circuit source and destination, and when created results in the creation of high-order pass-through circuits on each node in the path. For any of the listed topology combinations except SNCP to SNCP, one tunnel is required; an SNCP-to-SNCP topology requires two tunnels.

7.18.2 Manual Low-order Circuit Creation

When you must route low-order traffic from an ONS 15454 SDH SNCP across an ONS 15600 SDH hub to a line-protected domain, CTC cannot automatically create the end-to-end low-order circuit. The configuration requires three TL1-like circuits. You must create all three circuits separately:

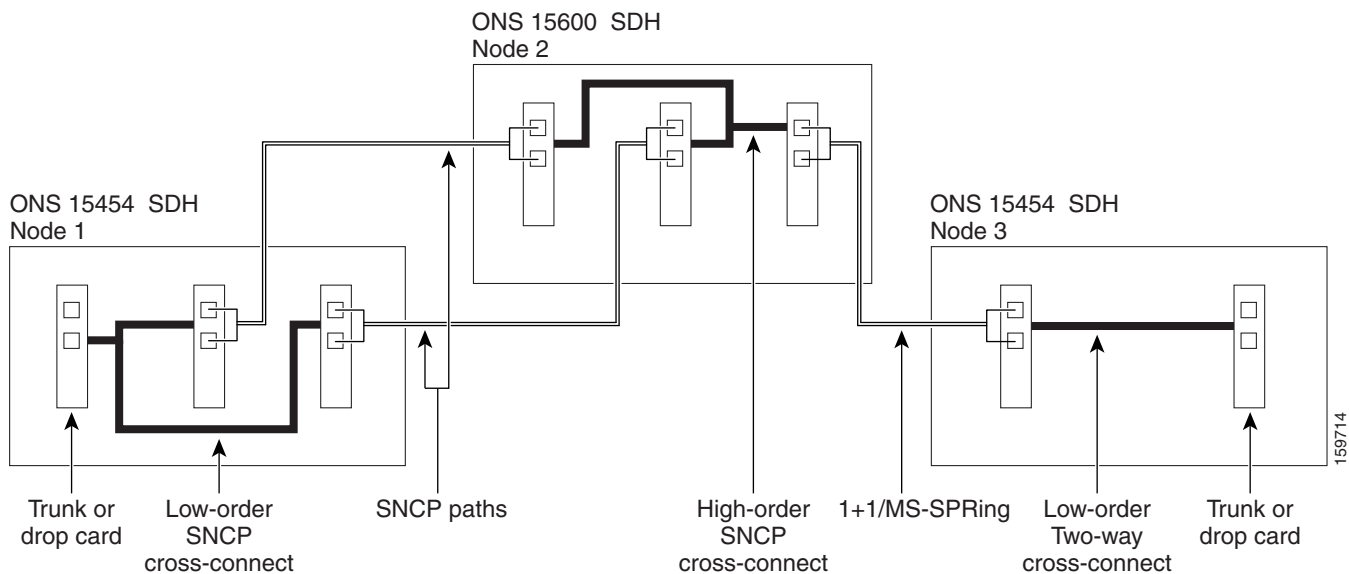
1. Low-order open-ended circuit on the ONS 15454 SDH SNCP
2. STS open-ended circuit on the ONS 15600 SDH hub node
3. Low-order two-way circuit on the ONS 15454 SDH line-protected node

The three TL1-like circuits must physically align for CTC to automatically merge the circuits.

[Figure 7-14](#) is an example of an ONS 15600 SDH node bridging an ONS 15454 SDH SNCP and an ONS 15454 SDH line-protected domain. In [Figure 7-14](#), the following circuits are present:

- An open-ended low-order circuit exists on Node 1. Slot 1 is the source, and Slots 2 and 3 are the destinations.
- An open-ended high-order circuit exists on Node 2 (the hub ONS 15600 SDH node). Slot 1 and Slot 2 are the sources, and Slot 3 is the destination.
- A two-way low-order circuit exists on Node 3 between Slot 1 and Slot 2.

Figure 7-14 Low-order Traffic Routing over an ONS 15600 SDH Hub Node



For information on creating an open-ended circuit on the ONS 15600 SDH to bridge protection domains, refer to the “Turn Up Network” chapter in the *Cisco ONS 15600 SDH Procedure Guide*.

7.19 High-Order VC3 Traffic Routing

High-order VC3 circuits are supported on the ONS 15600 SDH for connecting an ONS 15600 SDH node to another ONS 15600 SDH node, or for connecting an ONS 15600 SDH node to third-party SDH network elements. You can create VC3 circuits that use STM-1, STM-4, STM-16, and STM-64 interfaces.

