CHAPTER 10

Circuits and Tunnels

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

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.

Chapter topics include:

10.1 Overview

On an ONS 15454, you can create unidirectional and bidirectional circuits. For path protection circuits, you can create revertive or nonrevertive circuits. Circuits are routed automatically or you can manually route them. With the autorange feature, you do not need to individually build multiple circuits of the same type; CTC can create additional sequential circuits if you specify the number of circuits you need and build the first circuit.

You can provision circuits at any of the following points:

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

10.2 Circuit Properties

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

- Name—The name of the circuit. The circuit name can be manually assigned or automatically generated.
- Type—The circuit types are STS (STS circuit), VT (VT circuit), VTT (VT tunnel), VAP (VT aggregation point), OCHNC (dense wavelength division multiplexing [DWDM] optical channel network connection), or STS-V (STS virtual concatenated [VCAT] circuit).
- Size—The circuit size. VT circuits are 1.5. STS circuit sizes are 1, 3c, 6c, 9c, 12c, 24c, 36c, 48c, 192c. OCHNC sizes are Equipped non specific, Multi-rate, 2.5 Gbps No FEC (forward error correction), 2.5 Gbps FEC, 10 Gbps No FEC, and 10 Gbps FEC. VCAT circuits are STS-1-2v, STS-3c-2v, and STS-12c-2v.
- OCHNC Wlen—For OCHNCs, the wavelength provisioned for the optical channel network connection.
- Direction—The circuit direction, either two-way or one-way.
- OCHNC Dir—For OCHNCs, the direction of the optical channel network connection, either east to west or west to east.
- Protection—The type of circuit protection. See the "10.2.3 Circuit Protection Types" section on page 10-6 for a list of protection types.
- Status—The circuit status. See the "10.2.1 Circuit Status" section on page 10-3.
- Source—The circuit source in the format: node/slot/port "port name"/STS/VT. (The port name appears in quotes.) Node and slot always appear; port "port name"/STS/VT might appear, depending on the source card, circuit type, and whether a name is assigned to the port. If the circuit size is a concatenated size (3c, 6c, 12c, etc.), STSs used in the circuit are indicated by an ellipsis, for example, S7..9, (STSs 7, 8, and 9) or S10..12 (STS 10, 11, and 12).

- Destination—The circuit destination in same format (node/slot/port “port name”/STS/VT) as the circuit source.
- # of VLANS—The number of VLANS used by an Ethernet circuit.
- # of Spans—The number of inter-node links that constitute the circuit. Right-clicking the column displays a shortcut menu from which you can choose to show or hide circuit span detail.
- State—The circuit state. See the “10.2.2 Circuit States” section on page 10-5.

Figure 10-1 ONS 15454 Circuit Window in Network View

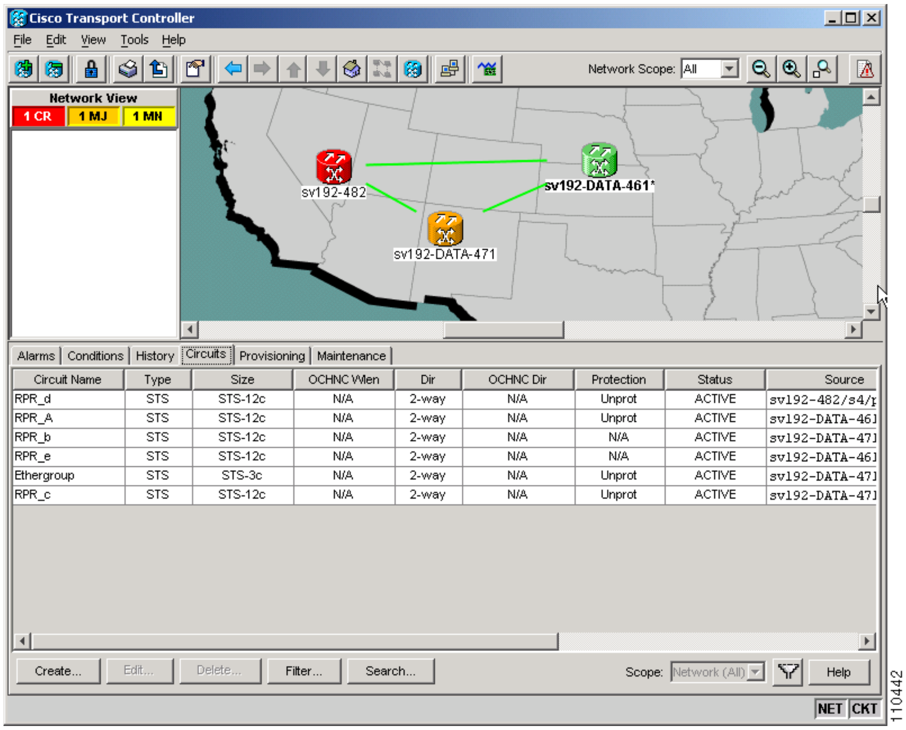

Circuit Name	Type	Size	OCHNC Wlen	Dir	OCHNC Dir	Protection	Status	Source
RPR_d	STS	STS-12c	N/A	2-way	N/A	Unprot	ACTIVE	sv192-482/s4/p
RPR_A	STS	STS-12c	N/A	2-way	N/A	Unprot	ACTIVE	sv192-DATA-461
RPR_b	STS	STS-12c	N/A	2-way	N/A	N/A	ACTIVE	sv192-DATA-471
RPR_e	STS	STS-12c	N/A	2-way	N/A	N/A	ACTIVE	sv192-DATA-461
Ethergroup	STS	STS-3c	N/A	2-way	N/A	Unprot	ACTIVE	sv192-DATA-471
RPR_c	STS	STS-12c	N/A	2-way	N/A	Unprot	ACTIVE	sv192-DATA-471

10.2.1 Circuit Status

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

Table 10-1 ONS 15454 Circuit Status

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

Table 10-1 ONS 15454 Circuit Status (continued)

Status	Definition/Activity
INCOMPLETE	A CTC-created circuit is missing a cross-connect or network span, a complete path from source to destination(s) does not exist, or an Alarm Interface Panel (AIP) change occurred on one of the circuit nodes and the circuit is in need of repair. (AIPs store the node MAC address.) In CTC, circuits are represented using cross-connects and network spans. If a network span is missing from a circuit, the circuit status is INCOMPLETE. However, an INCOMPLETE status does not necessarily mean a circuit traffic failure has occurred, because traffic may flow on a protect path. Network spans are in one of two states: up or down. On CTC circuit and network maps, up spans appear as green lines, and down spans appear as gray lines. If a failure occurs on a network span during a CTC session, the span remains on the network map but its color changes to gray to indicate that the span is down. If you restart your CTC session while the failure is active, the new CTC session cannot discover the span and its span line does not appear on the network map. Subsequently, circuits routed on a network span that goes down appear as ACTIVE during the current CTC session, but appear as INCOMPLETE to users who log in after the span failure. INCOMPLETE status does not apply to OCHNC circuit types.
UPGRADABLE	A TL1-created circuit or a TL1-like, CTC-created circuit is complete and has upgradable cross-connects. A complete path from source to destination(s) exists. The circuit can be upgraded. This status does not apply to OCHNC circuit types.
INCOMPLETE_UPGRADABLE	A TL1-created circuit or a TL1-like, CTC-created circuit with upgradable cross-connects is missing a cross-connect or circuit span (network link), and a complete path from source to destination(s) does not exist. The circuit cannot be upgraded until missing components are in place. This status does not apply to OCHNC circuit types.
NOT_UPGRADABLE	A TL1-created circuit or a TL1-like, CTC-created circuit is complete but has at least one non-upgradable cross-connect. UPSR_HEAD, UPSR_EN, UPSR_DC, and UPSR_DROP connections are not upgradable, so all path protection circuits created with TL1 are not upgradable. This status does not apply to OCHNC circuit types.
INCOMPLETE_NOT_UPGRADABLE	A TL1-created circuit or a TL1-like CTC-created circuit with one or more nonupgradable cross-connects is missing a cross-connect or circuit span (network link); a complete path from source to destination(s) does not exist. This status does not apply to OCHNC circuit types.

10.2.2 Circuit States

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

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

You can assign a state to circuits at two points:

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

Table 10-2 Circuit States

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

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

Table 10-3 Partial Circuit States

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

PARTIAL states can occur during automatic or manual transitions between states. OOS-AINS-PARTIAL appears if you assign OOS-AINS to a circuit with DS-1 or DS3XM cards as the source or destination. Some cross-connects transition to IS, while others are OOS-AINS. PARTIAL can appear during a manual transition caused by an abnormal event such as a CTC crash or communication error, or if one of the cross-connects could not be changed. Refer to the *Cisco ONS 15454 Troubleshooting Guide* for troubleshooting procedures.

Circuits do not use the soak timer for transitional states, but ports do. The soak period is the amount of time that the port remains in the OOS-AINS state after a signal is continuously received. When provisioned as OOS-AINS, the ONS 15454 monitors a circuit's cross-connects for an error-free signal. It changes the state of the circuit from OOS-AINS to IS or to AINS-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 DS-1 and DS-3 circuits using CTC are:

- When provisioning VT1.5 circuits and VT tunnels as OOS-AINS, the circuit state transitions to IS shortly after the circuits are created when the circuit source and destination ports are IS, OOS-AINS, or OOS-MT. The source and destination ports on the VT1.5 circuits remain in the OOS-AINS state until an alarm-free signal is received for the duration of the soak timer. When the soak timer expires, the VT1.5 source port and destination port states change to IS.
- When provisioning STS circuits as OOS-AINS, the circuit source and destination ports are OOS-AINS. As soon as an alarm-free signal is received the circuit state changes to IS and the source and destination ports remain OOS-AINS for the duration of the soak timer. After the port soak timer expires, STS source and destination ports change to IS.

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

10.2.3 Circuit Protection Types

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

Table 10-4 Circuit Protection Types

Protection Type	Description
—	Circuit protection is not applicable.
2F BLSR	The circuit is protected by a two-fiber bidirectional line switched ring (BLSR).
4F BLSR	The circuit is protected by a four-fiber BLSR.
BLSR	The circuit is protected by a both a two-fiber and a four-fiber BLSR.
Path Protection	The circuit is protected by a path protection.
Path Protection-DRI	The circuit is protected by a path protection dual ring interconnection
1+1	The circuit is protected by a 1+1 protection group.
Y-Cable	The circuit is protected by a transponder or muxponder card Y-cable protection group.
Protected	The circuit is protected by diverse SONET topologies, for example, a BLSR and a path protection, or a path protection and 1+1.

Table 10-4 Circuit Protection Types (continued)

Protection Type	Description
2F-PCA	The circuit is routed on a protection channel access (PCA) path on a two-fiber BLSR. PCA circuits are unprotected.
4F-PCA	The circuit is routed on a protection channel access path on a four-fiber BLSR. PCA circuits are unprotected.
PCA	The circuit is routed on a protection channel access path on both two-fiber and four-fiber BLSRs. PCA circuits are unprotected.
Unprot (black)	The circuit is not protected.
Unprot (red)	A circuit created as a fully-protected circuit is no longer protected due to a system change, such as a traffic switch.
Unknown	The circuit protection types appear in the Protection column only when all circuit components are known, that is, when the circuit status is ACTIVE or UPGRADABLE. If the circuit is in some other status, the protection type appears as "unknown."

10.2.4 Circuit Information in the Edit Circuit Window

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

- Circuit direction (unidirectional/bidirectional)
- The nodes, STSs, and VTs through which a circuit passes, including slots and port numbers
- The circuit source and destination points
- OSPF Area IDs
- Link protection (path protection, unprotected, BLSR, 1+1) and bandwidth (OC-N)

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

You can also view alarms and states 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

Figure 10-2 shows a bidirectional STS circuit routed on a path protection.

Figure 10-2 Path Protection Circuit Displayed on the Detailed Circuit Map

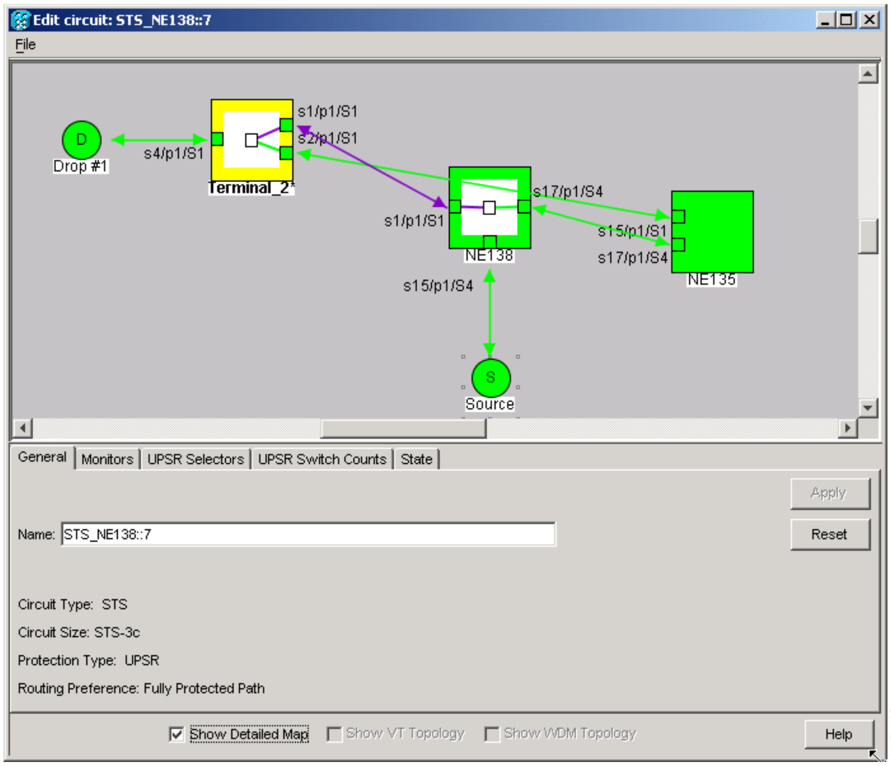

Name: STS_NE138::7

Circuit Type: STS

Circuit Size: STS-3c

Protection Type: UPSR

Routing Preference: Fully Protected Path

☑ Show Detailed Map □ Show VT Topology □ Show WDM Topology

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

Table 10-5 Port State Color Indicators

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

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

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

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

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

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

Figure 10-3 on page 10-10 shows an example of the information that can appear. From this example, you can determine:

- The circuit has one source and one destination.
- The circuit has three nodes in its route; the state of the most severe alarm can be determined by the color of the node icons. For example, yellow indicates that the most severe alarm is minor in severity.
- The STSs and ports that the circuit passes through from source to destination.
- The port states and severity of the most severe alarm on each port.
- A facility loopback exists on the port at one end of the circuit; a terminal loopback exists at the other end port.
- An automatic path trace exists on one STS end of the circuit; a manual path trace exists at the other STS end.
- The circuit is path protection-protected (by path selectors). One path selector has a Lockout, one has a Force switch, one has a Manual switch, and the others are free of external switching commands.
- The working path (green) flows from ptlm6-454a59-24/s6/p1/S1 to dv9-241/s6/p1/S1, and from dv9-241/s16/p1/S1 to tccp/s14/p1/vc3-3. The protect path (purple) is also visible.
- On ptlm6-454a59-24 and tccp, the working path is active; on dv9-241, the protect path is active.

From the example, you could:

- Display any port or node view.
- Edit the path trace states of any port that supports path trace.
- Change the path selector state of any path protection path selector.

Figure 10-3 Detailed Circuit Map Showing a Terminal Loopback

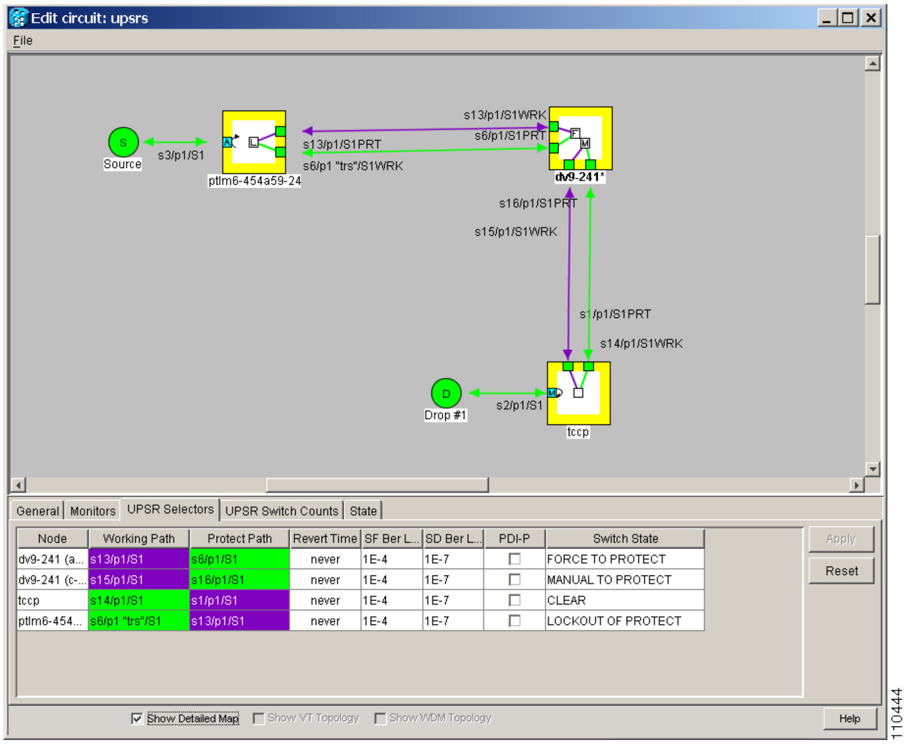

Node	Working Path	Protect Path	Revert Time	SF Ber L...	SD Ber L...	PDI-P	Switch State
dv9-241 (a...	s13/p1/S1	s6/p1/S1	never	1E-4	1E-7	□	FORCE TO PROTECT
dv9-241 (c-...	s15/p1/S1	s16/p1/S1	never	1E-4	1E-7	□	MANUAL TO PROTECT
tccp	s14/p1/S1	s1/p1/S1	never	1E-4	1E-7	□	CLEAR
ptlm6-454...	s6/p1 "trs"/S1	s13/p1/S1	never	1E-4	1E-7	□	LOCKOUT OF PROTECT

☑ Show Detailed Map ☐ Show VT Topology ☐ Show WDM Topology

10.3 Cross-Connect Card Bandwidth

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

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

The XCVT and XC10G cards perform VT1.5 multiplexing through 24 logical STS ports on the XCVT or XC10G VT matrix. Each logical STS port can carry 28 VT1.5s. Subsequently, the VT matrix has capacity for 672 VT1.5s terminations, or 336 VT1.5 circuits, because every circuit requires two terminations, one for ingress and one for egress. However, this capacity is only achievable if:

- Every STS port on the VT matrix carries 28 VT1.5s.
- The node is in a BLSR.

For example, if you create a VT1.5 circuit from STS-1 on a drop card and a second VT1.5 circuit from STS-2, two VT matrix STS ports are used, as shown in Figure 10-4. If you create a second VT1.5 circuit from the same STS port on the drop card, no additional logical STS ports are used on the VT matrix. However, if the next VT1.5 circuit originates on a different STS, a second STS port on the VT matrix is used, as shown in Figure 10-5. If you continued to create VT1.5 circuits on a different EC-1 STSs and mapped each to an unused outbound STS, the VT matrix capacity would be reached after you created 12 VT1.5 circuits.

Figure 10-4 One VT1.5 Circuit on One STS

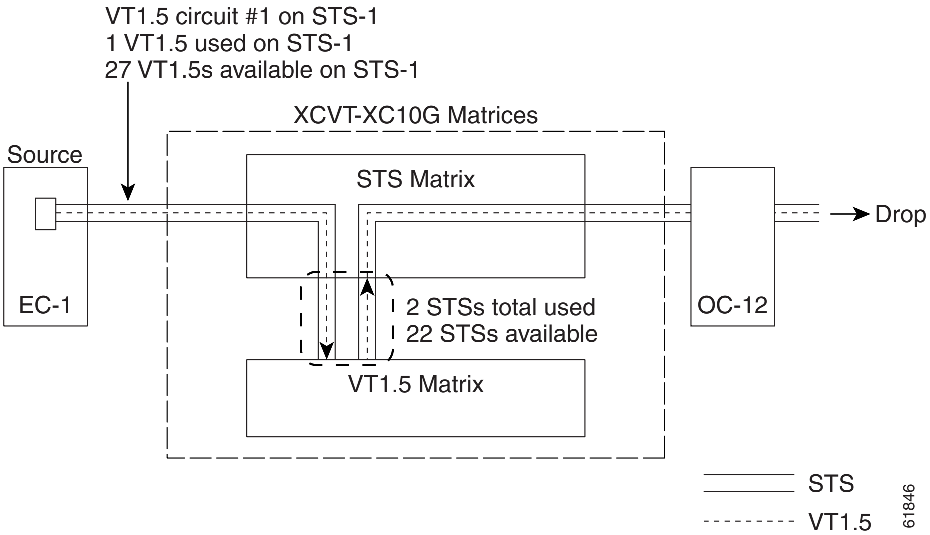

Figure 10-5 Two VT1.5 Circuits in a BLSR

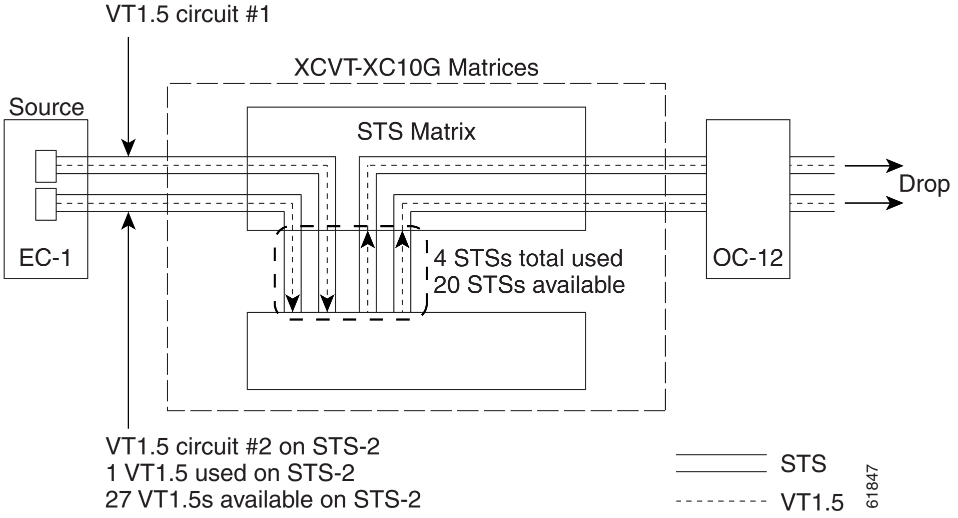

Note Circuits with DS1-14 and DS1N-14 circuit sources or destinations use one STS port on the VT matrix. Because you can only create 14 VT1.5 circuits from the DS-1 cards, 14 VT1.4s are unused on the VT matrix.

VT matrix capacity is also affected by SONET protection topology and node position within the circuit path. Matrix usage is slightly higher for path protection and 1+1 nodes than BLSR nodes. Circuits use two VT matrix ports at pass-through nodes if VT tunnels and aggregation points are not used. If the circuit is routed on a VT tunnel or an aggregation point, no VT matrix resources are used. Table 10-6 shows basic STS port usage rates for VT 1.5 circuits.

Table 10-6 VT Matrix Port Usage for One VT1.5 Circuit

Node Type	No Protection	BLSR	Path Protection	1+1
Circuit source or destination node	2	2	3	3
Circuit pass-through node without VT tunnel	2	2	2	4
Circuit pass-through node with VT tunnel	0	0	0	0

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

- STS-1 Matrix—The percent of STS matrix resources that are used. 288 STSs are available on XC and XCVT cards; 1152 are available on XC10G cards.
- VT Matrix Ports—The percent of the VT matrix ports (logical STS ports) that are used. No ports are available on XC cards; 24 are available on XCVT and XC10G cards. The VT Port Matrix Detail shows the percent of each VT matrix port that is used.
- VT Matrix—The percent of the total VT matrix terminations that are used. There are 672 terminations, which is the number of logical STS VT matrix ports (24) multiplied by the number of VT1.5s per port (28).

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

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

10.4 DCC Tunnels

SONET provides four DCCs for network element operations, administration, maintenance, and provisioning: one on the SONET Section layer (DCC1) and three on the SONET Line layer (DCC2, DCC3, and DCC4). The ONS 15454 uses the section DCC for ONS 15454 management and provisioning. A section DCC (SDCC) and line DCC (LDCC) each provide 192 Kbps of bandwidth per channel. The aggregate bandwidth of the three LDCCs is 576 Kbps. When multiple DCC channels exist between two neighboring nodes, the ONS 15454 balances traffic over the existing DCC channels using a load balancing algorithm. This algorithm chooses a DCC for packet transport by considering packet size and DCC utilization.

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

10.4.1 Traditional DCC Tunnels

In traditional DCC tunnels, you can use the three line DCCs and the section DCC (when not used for ONS 15454 DCC terminations). A traditional DCC tunnel endpoint is defined by slot, port, and DCC, where DCC can be either the section DCC or one of the line DCCs. You can link line DCCs to line DCCs and link section DCCs to section DCCs. You can also link a section DCC to an line DCC, and a line DCC to a section DCC. To create a DCC tunnel, you connect the tunnel endpoints from one ONS 15454 optical port to another. Cisco recommends a maximum of 84 DCC tunnel connections for an ONS 15454. Table 10-7 shows the DCC tunnels that you can create using different OC-N cards.

Table 10-7 DCC Tunnels

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

Figure 10-6 shows a DCC tunnel example. Third-party equipment is connected to OC-3 cards at Node 1/Slot 3/Port 1 and Node 3/Slot 3/Port 1. Each ONS 15454 node is connected by OC-48 trunk (span) cards. In the example, three tunnel connections are created, one at Node 1 (OC-3 to OC-48), one at Node 2 (OC-48 to OC-48), and one at Node 3 (OC-48 to OC-3).

Figure 10-6 Traditional DCC Tunnel

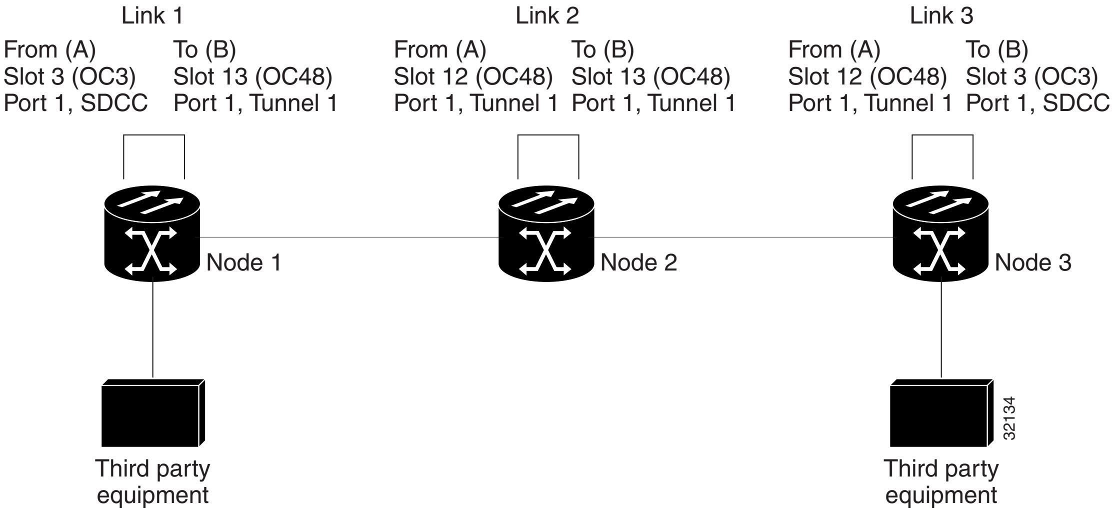

When you create DCC tunnels, keep the following guidelines in mind:

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

10.4.2 IP-Encapsulated Tunnels

An IP-encapsulated tunnel puts a section 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 SDCC bandwidth.

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

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

10.5 Multiple Destinations for Unidirectional Circuits

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

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

10.6 Monitor Circuits

Monitor circuits are secondary circuits that monitor traffic on primary bidirectional circuits. Figure 10-7 shows an example of a monitor circuit. At Node 1, a VT1.5 is dropped from Port 1 of an EC1-12 card. To monitor the VT1.5 traffic, plug test equipment into Port 2 of the EC1-12 card and provision a monitor circuit to Port 2. Circuit monitors are one-way. The monitor circuit in Figure 10-7 monitors VT1.5 traffic received by Port 1 of the EC1-12 card.

Figure 10-7 VT1.5 Monitor Circuit Received at an EC1-12 Port

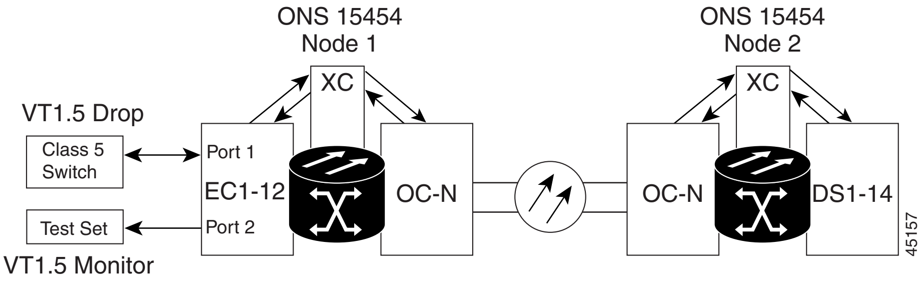

Note Monitor circuits cannot be used with Ethernet circuits.

10.7 Path Protection Circuits

Use the Edit Circuits window to change path protection selectors and switch protection paths (Figure 10-8). In this window, you can:

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

Figure 10-8 Editing Path Protection Selectors

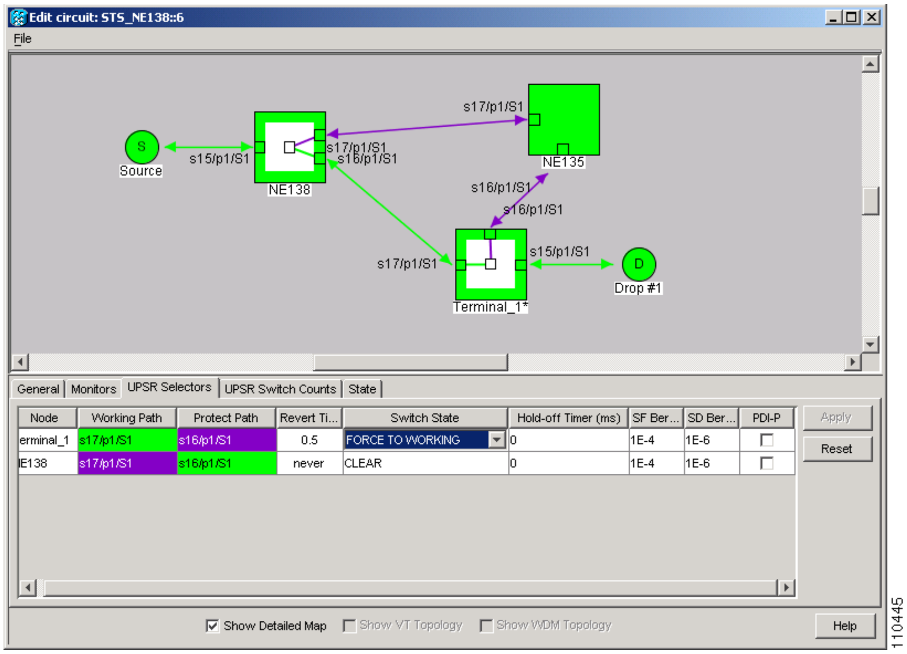

Node	Working Path	Protect Path	Revert Ti...	Switch State	Hold-off Timer (ms)	SF Ber...	SD Ber...	PDI-P
erminal_1	s17/p1/S1	s16/p1/S1	0.5	FORCE TO WORKING	0	1E-4	1E-6	□
IE138	s17/p1/S1	s16/p1/S1	never	CLEAR	0	1E-4	1E-6	□

10.7.1 Open-Ended Path Protection Circuits

If ONS 15454s are connected to a third-party network, you can create an open-ended path protection circuit to route a circuit through it. To do this, you create three circuits. One circuit is created on the source ONS 15454 network. This circuit has one source and two destinations, one at each ONS 15454 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 15454s. That circuit routes the two circuit signals across the network to ONS 15454s 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 path protection circuit.

10.7.2 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 10-9). The working path will always be the shortest path. If a fault occurs, both the working and protection fibers are not affected. This feature only applies to bidirectional path protection circuits. The go-and-return option appears on the Circuit Attributes panel of the Circuit Creation wizard.

Figure 10-9 Path Protection Go-and-Return Routing

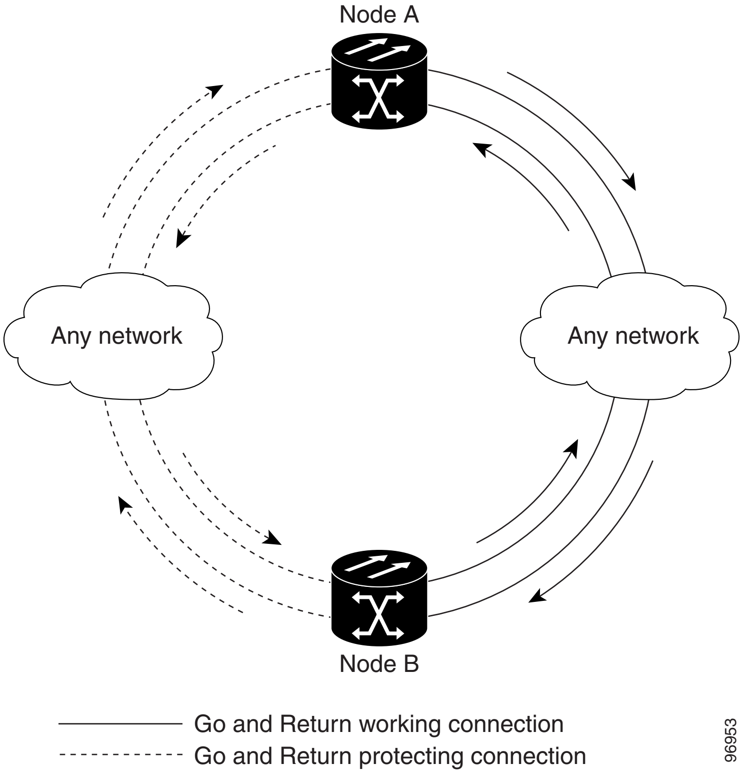
Node A
Any network
Any network
Node B
Go and Return working connection
Go and Return protecting connection
96953

10.8 BLSR Protection Channel Access Circuits

You can provision circuits to carry traffic on BLSR protection channels when conditions are fault-free. Traffic routed on BLSR PCA circuits, called extra traffic, has lower priority than the traffic on the working channels and has no means for protection. During ring or span switches, PCA circuits are preempted and squelched. For example, in a two-fiber OC-48 BLSR, STSs 25-48 can carry extra traffic when no ring switches are active, but PCA circuits on these STSs are preempted when a ring switch occurs. When the conditions that caused the ring switch are remedied and the ring switch is removed, PCA circuits are restored. If the BLSR is provisioned as revertive, this occurs automatically after the fault conditions are cleared and the reversion timer has expired.

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

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

10.9 Path Trace

The SONET J1 Path Trace is a repeated, fixed-length string comprised of 64 consecutive J1 bytes. You can use the string to monitor interruptions or changes to circuit traffic. Table 10-8 shows the ONS 15454 cards that support path trace. DS-1 and DS-3 cards can transmit and receive the J1 field, while the EC-1, OC-3, OC-48AS, and OC-192 can only receive the J1 bytes. Cards that are not listed in the table do not support the J1 byte.

Table 10-8 ONS 15454 Cards Capable of Path Trace

J1 Function	Cards
Transmit and Receive	DS1-14
	DS1N-14
	DS3-12E
	DS3i-N-12
	DS3N-12E
	DS3XM-6
	G-Series
	ML-Series
Receive Only	EC1-12
	OC3 IR 4 1310
	OC12/STM4-4
	OC48 IR/STM16 SH AS 1310
	OC48 LR/STM16 LH AS 1550
	OC192 LR/STM64 LH 1550

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

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

10.10 Path Signal Label, C2 Byte

One of the overhead bytes in the SONET frame is the C2 Byte. The SONET standard defines the C2 byte as the path signal label. The purpose of this byte is to communicate the payload type being encapsulated by the STS path overhead (POH). The C2 byte functions similarly to EtherType and Logical Link Control

(LLC)/Subnetwork Access Protocol (SNAP) header fields on an Ethernet network; it allows a single interface to transport multiple payload types simultaneously. C2 byte hex values are provided in Table 10-9.

Table 10-9 STS Path Signal Label Assignments for Signals

Hex Code	Content of the STS SPE
0x00	Unequipped
0x01	Equipped - nonspecific payload
0x02	Virtual Tributary (VT) structured STS-1 (DS-1)
0x03	Locked VT mode
0x04	Asynchronous mapping for DS-3
0x12	Asynchronous mapping for DS4NA
0x13	Mapping for Asynchronous Transfer Mode (ATM)
0x14	Mapping for distributed queue dual bus (DQDB)
0x15	Asynchronous mapping for fiber distributed data interface (FDDI)
0x16	High level data link control (HDLC) over SONET mapping
0xFD	Reserved
0xFE	0.181 Test signal (TSS1 to TSS3) mapping SDH network
0xFF	Alarm indication signal, path (AIS-P)

If a circuit is provisioned using a terminating card, the terminating card provides the C2 byte. A VT circuit is terminated at the XCVT or XC-10G card, which generates the C2 byte (0x02) downstream to the STS terminating cards. The XCVT or XC10G card generates the C2 value (0x02) to the DS1 or DS3XM terminating card. If an optical circuit is created with no terminating cards, the test equipment must supply the path overhead in terminating mode. If the test equipment is in "pass through mode," the C2 values usually change rapidly between 0x00 and 0xFF. Adding a terminating card to an optical circuit usually fixes a circuit having C2 byte problems. Table 10-10 lists label assignments for signals with payload defects.

Table 10-10 STS Path Signal Label Assignments for Signals with Payload Defects

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

Table 10-10 STS Path Signal Label Assignments for Signals with Payload Defects (continued)

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

10.11 Automatic Circuit Routing

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

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

- Circuit routing tries to use the shortest path within the user-specified or network-specified constraints. VT tunnels are preferable for VT circuits because VT tunnels are considered shortcuts when CTC calculates a circuit path in path-protected mesh networks.
- If you do not choose fully path protected during circuit creation, circuits can still contain protected segments. Because circuit routing always selects the shortest path, one or more links and/or segments can have some protection. CTC does not look at link protection while computing a path for unprotected circuits.

- Circuit routing does not use links that are down. If you want all links to be considered for routing, do not create circuits when a link is down.
- Circuit routing computes the shortest path when you add a new drop to an existing circuit. It tries to find the shortest path from the new drop to any nodes on the existing circuit.
- If the network has a mixture of VT-capable nodes and VT-incapable nodes, CTC may automatically create a VT tunnel. Otherwise, CTC asks you whether a VT tunnel is needed.
- You cannot create protected VT circuits between path protection configurations and BLSRs if an XC card is installed on the node shared by the two topologies. To create protected circuits between topologies, install an XCVT or XC10G cross-connect card on the shared node.

10.11.1 Bandwidth Allocation and Routing

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

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

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

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

10.11.2 Secondary Sources and Destination

CTC supports secondary circuit sources and destinations (drops). Secondary sources and destinations can be created to connect two third-party networks, as shown in Figure 10-10. Traffic is protected while it goes through a network of ONS 15454s.

Figure 10-10 Secondary Sources and Destinations

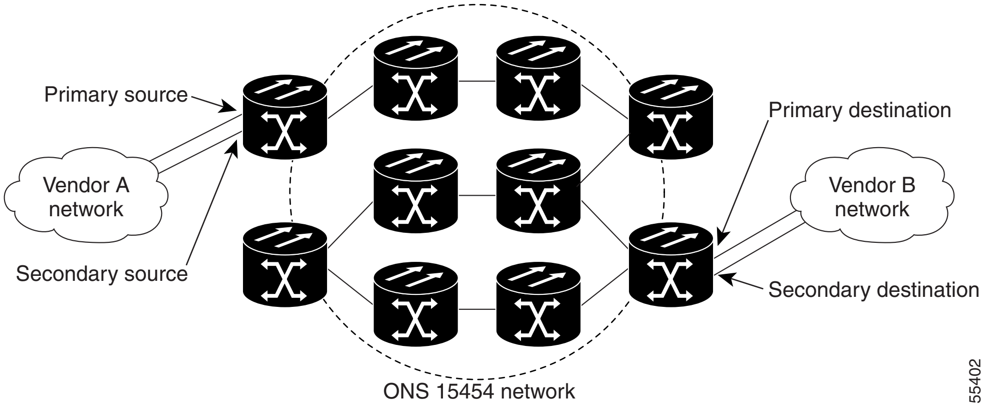

Several rules apply to secondary sources and destinations:

- CTC does not allow a secondary destination for unidirectional circuits because you can always specify additional destinations after you create the circuit.
- The sources and destinations cannot be DS-3, DS3XM, or DS-1-based STS-1s or VT1.5s.
- Secondary sources and destinations are permitted only for regular STS/VT1.5 connections (not for VT tunnels and multicard EtherSwitch circuits).
- For point-to-point (straight) Ethernet circuits, only SONET STS endpoints can be specified as multiple sources or destinations.

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

For unidirectional circuits, a path protection drop-and-continue connection is created at the source node.

10.12 Manual Circuit Routing

Routing circuits manually allows you to:

- Choose a specific path, not necessarily the shortest path.
- Choose a specific STS/VT1.5 on each link along the route.
- Create a shared packet ring for multicard EtherSwitch circuits.
- Choose a protected path for multicard EtherSwitch circuits, allowing virtual path protection segments.

CTC imposes the following rules on manual routes:

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

Figure 10-11 Alternate Paths for Virtual Path Protection Segments

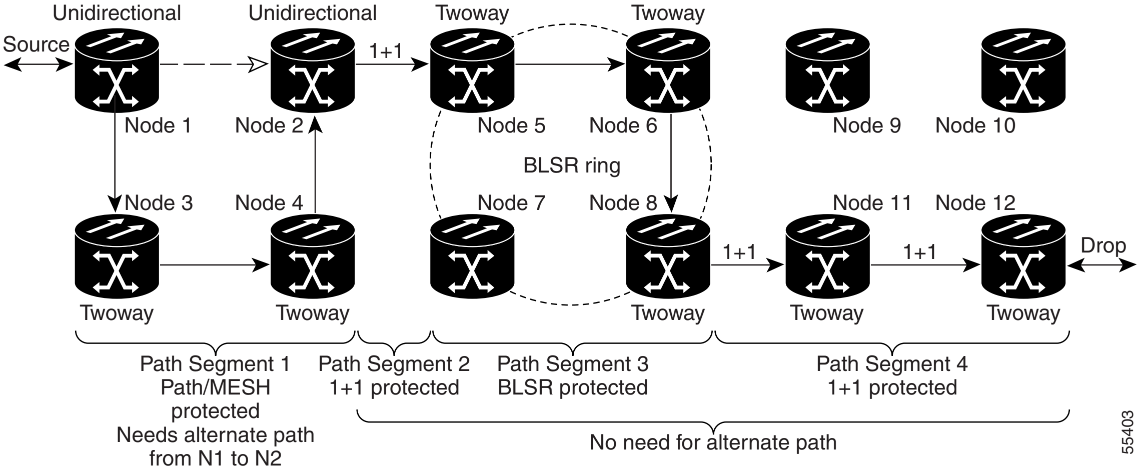

- For multicard EtherSwitch circuits, the fully path protected option is ignored.
- For a node that has a path protection selector based on the links chosen, the input links to the path protection selectors cannot be 1+1 or BLSR protected (Figure 10-12). The same rule applies at the path protection bridge.

Figure 10-12 Mixing 1+1 or BLSR Protected Links With a Path Protection

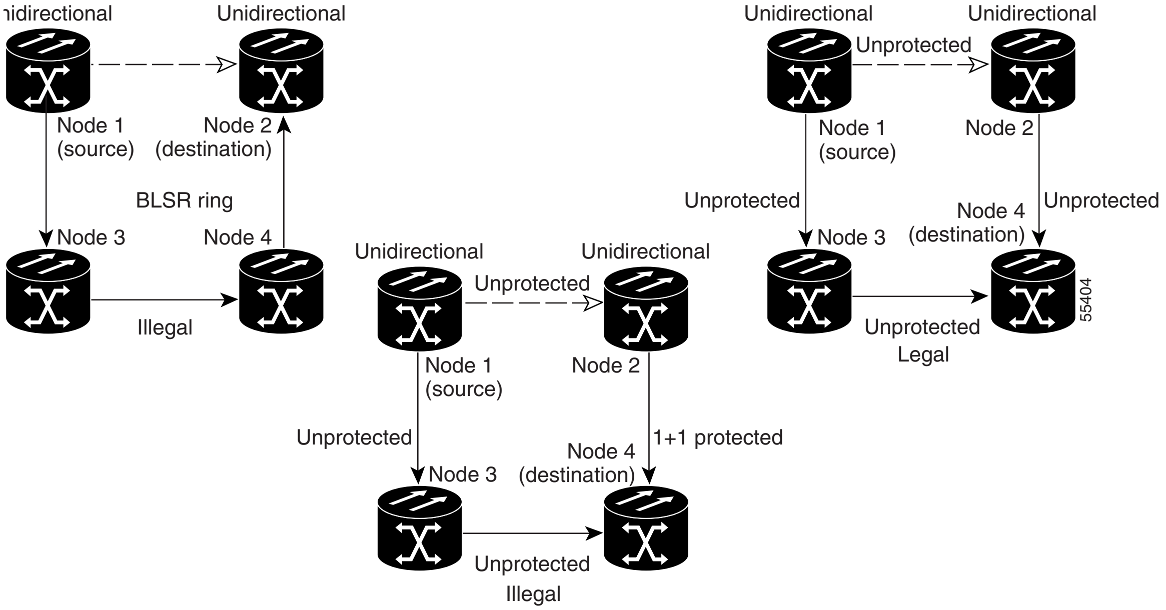

- Choose the links of multicard EtherSwitch circuits in a shared packet ring to route from source to destination back to source (Figure 10-13). Otherwise, a route (set of links) chosen with loops is invalid.

Figure 10-13 Ethernet Shared Packet Ring Routing

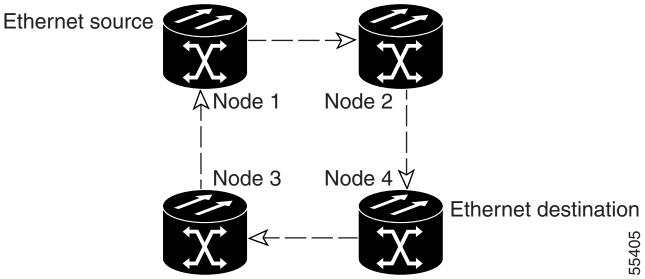

- Multicard EtherSwitch circuits can have virtual path protection segments if the source or destination is not in the path protection domain. This restriction also applies after circuit creation; therefore, if you create a circuit with path protection segments, Ethernet destinations cannot exist anywhere on the path protection segment (Figure 10-14).

Figure 10-14 Ethernet and Path Protection

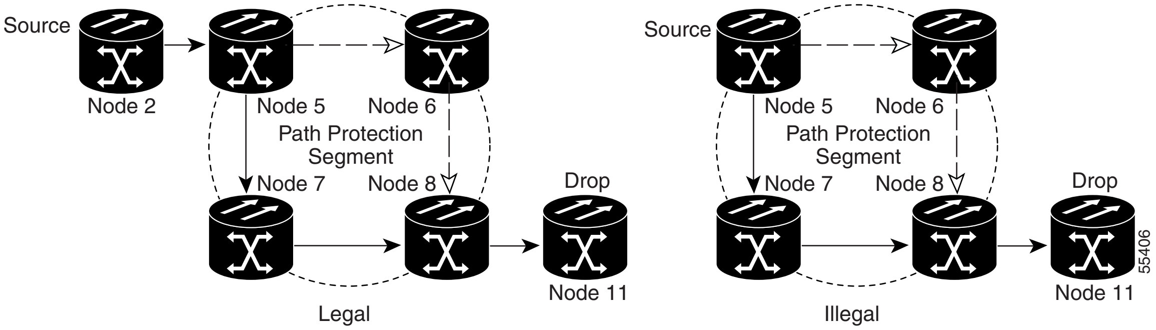

- VT tunnels cannot be the endpoint of a path protection segment. A path protection segment endpoint is where the path protection selector resides.

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

Table 10-11 through Table 10-14 on page 10-25 summarize the available node connections. Any other combination is invalid and generates an error.

Table 10-11 Bidirectional STS/VT/Regular Multicard EtherSwitch/Point-to-Point (Straight) Ethernet Circuits

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

Table 10-11 Bidirectional STS/VT/Regular Multicard EtherSwitch/Point-to-Point (Straight) Ethernet Circuits (continued)

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

Table 10-12 Unidirectional STS/VT Circuit

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Destinations
One way	1	1	—	—
Path Protection headend	1	2	—	—
Path Protection headend	—	2	1	—
Path Protection drop and continue	2	—	—	1+

Table 10-13 Multicard Group Ethernet Shared Packet Ring Circuit

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Destinations
At intermediate nodes only				
Double path protection	2	2	—	—
Two way	1	1	—	—
At source or destination nodes only				
Ethernet	1	1	—	—

Table 10-14 Bidirectional VT Tunnels

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Destinations
At intermediate nodes only				
Path Protection	2	1	—	—

Table 10-14 Bidirectional VT Tunnels (continued)

Connection Type	Number of Inbound Links	Number of Outbound Links	Number of Sources	Number of Destinations
Path Protection	1	2	—	—
Double path protection	2	2	—	—
Two way	1	1	—	—
At source nodes only				
VT tunnel endpoint	—	1	—	—
At destination nodes only				
VT tunnel endpoint	1	—	—	—

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

10.13 Constraint-Based Circuit Routing

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

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

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

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

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

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

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

10.14 Virtual Concatenated Circuits

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

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

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

The automatic and manual routing selection applies to the entire VCAT circuit, that is, all members are manually or automatically routed. In Software R4.6, bidirectional VCAT circuits are symmetric, which means that the same number of members travel in each direction. Software R4.6 supports common fiber routing, where all VCAT members travel on the same fibers, thus eliminating delay between members. Figure 10-15 shows an example of common fiber routing.

Figure 10-15 VCAT on Common Fiber

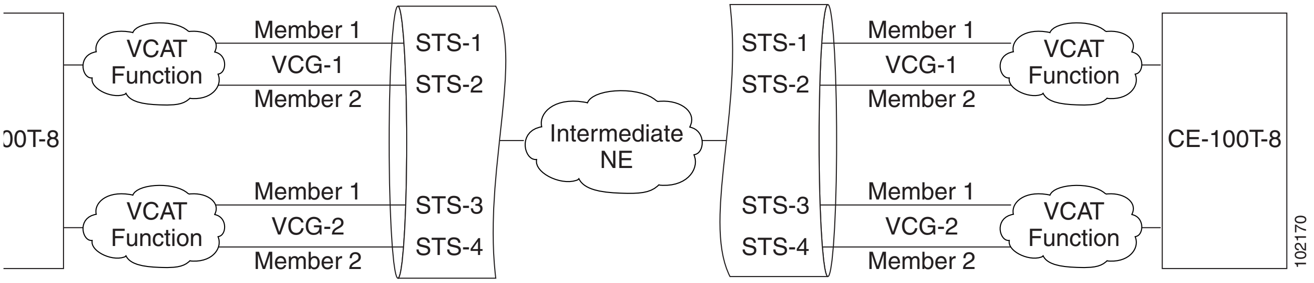

The Software–Link Capacity Adjustment Scheme (Sw-LCAS) uses legacy SONET failure indicators like the AIS-P and RDI-P to detect member failure. Sw-LCAS removes the failed member from the VCAT circuit for the duration of the failure, leaving the remaining members to carry the traffic. When the failure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

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