

DLSw+ Ethernet Redundancy Feature

This chapter describes how the DLSw+ Ethernet redundancy feature works, how to configure it, sample configurations, and configuration verification tips.

DLSw+ provides redundancy and load balancing when the end systems are connected over any media that support SRB: Token Ring, FDDI, Token Ring LAN Emulation (LANE), and Token Ring InterSwitch Link (ISL). When the end systems are connected to Ethernet, however, DLSw+ has design limitations in providing redundancy and load balancing. See Appendix C, “Ethernet DLSw+ Redundancy,” for a complete discussion on network design issues in a DLSw+ environment with Ethernet-attached end systems.

The DLSw+ Ethernet redundancy feature, introduced in Cisco IOS Release 12.0(5)T, provides redundancy and load balancing between multiple DLSw+ peers in an Ethernet environment. The feature also enables DLSw+ to support multiple DLSw+ routers on the same transparent bridged domain that can reach the same MAC address in a switched environment.

DLSw+ Ethernet Redundancy Feature Description

The DLSw+ Ethernet redundancy feature provides redundancy and load balancing between multiple DLSw+ peers in an Ethernet environment. It enables DLSw+ to support parallel paths between two points in an Ethernet environment, ensuring resiliency in the case of a router failure and providing load balancing for traffic load. This feature changes the way DLSw+ devices:

- Establish circuits in a transparent bridged domain
- Populate the SMAC field in explorer packets in a transparent bridged environment
- Support switch devices in an Ethernet environment

Establish Circuits in a Transparent Bridged Domain


The DLSw+ Ethernet redundancy feature alters the way DLSw+ routers establish circuits so that they recognize when they receive a frame from another router located on its transparent bridged domain. Before this feature, the devices connected to the same transparent bridged domain could not determine whether the frames they received were locally sourced, or whether they originated from another DLSw+ device within the same transparent bridged domain.

With the DLSw+ Ethernet redundancy feature enabled, the DLSw+ devices designate a master router in a transparent bridged domain. All devices on a transparent bridged domain advertise their presence to a multicast MAC address. One of the peers is elected as the master router. This master router maintains a database of all circuits being handled by the DLSw+ devices within its domain. Each device on the transparent bridged maintains

an LLC2 session with the master router and asks the master router for permission before starting or accepting a new DLSw+ circuit. Because the master router keeps a database of the circuits being handled, it prevents duplicate circuits from being created for the same SNA session.

In Figure 12-1, DLSw+ Routers A, B, and C are on the same transparent bridged domain. Router B is configured to be the master router.

Figure 12-1 Sample Diagram of DLSw+ with Ethernet Redundancy



In Figure 12-1, the following sequence occurs:

- End Station X sends a SABME to Y (which is seen by all three routers) to begin a NetBIOS session to End Station Y. Assume Routers A, B and C already have reachability information for End Station Y in their remote reachability cache at the time End Station X sends a SABME.
- Router A and Router C indicate they want the circuit by sending IWANTIT frames to Master Router B. IWANTIT frames are sent only in response to frames that start a circuit (SABME, XID); they are not sent in response to an explorer.
- Master Router B can also take the circuit. It waits, however, a designated amount of time before deciding which router gets the circuit in order to receive all qualified recipients. Master Router B bases its decision on the load information that is included in the IWANTIT frame of each router. In this particular case, we will assume that Master Router B decides that Router A should have the circuit and sends Router A an UGOTIT primitive frame.

- Master Router B also denies permission to Router C by sending it a CIRCUI_TAKEN (CKT_TKN) primitive frame. Master Router B then updates its own cache reachability tables, indicating that the circuit is taken.
- Router A sends a CANUREACH_cs to its remote peer to establish a circuit when it receives permission from the master router.

When the circuit disconnects, Router A notifies Master Router B by sending it a CIRCUI_GONE (CKT_GONE) primitive. The master router then forwards the CKT_GONE primitive to the other devices on the LAN and removes the circuit from its CKT_TKN database. The only time a master router deletes a record is when it is notified by the device to which it granted the circuit or when it loses its LLC2 session with that peer, for example, if there is a device failure.

Populate the SMAC Field in Explorer Packets in a Transparent Bridged Environment

The DLSw+ Ethernet redundancy feature changes the way a DLSw+ router replaces the SMAC of an explorer packet. Normally, when DLSw+ devices receive a TEST frame, they update their local or remote cache with the SMAC based on whether the packet came from its local LAN or off a WAN. In transparent bridged domains, this situation can create unreliable reachability information. With the DLSw+ Ethernet redundancy feature enabled, the SMAC of an explorer packet sent on the LAN is replaced by the DLSw+ router's own MAC address. When another router on the transparent bridged domain receives the explorer, it recognizes that the SMAC belongs to a DLSw+ router on its own LAN. (The routers in a transparent bridged domain learn of each other's MAC address during the master election process.) Therefore, it does not update its local reachability cache, and it does not forward the explorer over any of its peer connections, thereby enabling more reliable local cache reachability information and decreasing the chance for looping explorers.

Referring to Figure 12-1 again, DLSw+ Routers A, B, and C are on the same transparent bridged domain. The following sequence occurs:

- End Station Y sends an explorer looking for X. The remote peer sends a CANUREACH_ex to A, B and C across the WAN.
- Router C populates its REMOTE reachability cache with Y.
- Router C internally tracks that it is SEARCHING for device X and that when it finds device X, it must inform device Y (from the peer path from which the original CANUREACH_ex came).
- Router C transmits an explorer on the LAN because it does not have any reachability information on X. While doing so, it substitutes Y with its own MAC address in the SMAC field of the TEST frame and then it sends the TEST frame on the LAN.
- Routers A and B see the explorer from the LAN, but because they recognize Router C's MAC address in the SMAC field, they do not update their local reachability cache and do not forward the explorer over any of their DLSw+ peer connections. (Routers A, B and C have learned of each others MAC address during the master election process.)
- End Station X recognizes its MAC address in the TEST poll frame and responds to the SMAC, which is Router C.
- Routers A and B recognize this MAC address and do not act on the frame. The frame reaches Router C, which recognizes the frame as a response to its test poll.
- Router C updates its LOCAL cache with X and remembers that End Station Y was the original device searching for X. It replaces the MAC address with the original SMAC before sending an ICANREACH_ex reply back to the peer that originally sent the CANUREACH_ex.
- Routers A and B similarly respond to the remote peer's CANUREACH_ex by sending a TEST frame and substituting Y with their own MAC addresses.

Configuring the Ethernet Redundancy Feature

To configure the Ethernet redundancy feature, perform the following steps:


- Step 1. Issue the `dls transparent redundancy-enable` interface command on the Ethernet interface of all the DLSw+ devices located on the same transparent bridged domain. One of the options with this command is *multicast-mac-address*. *Configure the multicast-mac-address* option with the same MAC address on all DLSw+ devices configured for Ethernet redundancy. All the DLSw+ devices on the transparent bridged domain advertise their presence to this MAC address. In Figure 12-2, DLSw+ Routers A, B, and C are on the same transparent bridged domain and they all advertise their presence to multicast MAC address 9999.9999.9999.

When the routers that are within a transparent bridged domain learn each other's MAC address, all routers in the transparent bridged domain compete and elect a master based on the router with the lowest MAC address (that is, if it is not already configured based on the master-priority value.) If two masters are configured with an equal master priority setting, the router with the lowest MAC address is elected. (See the *Bridging and IBM Networking Command Reference* for command details.)

- Step 2. (Optional) Set the timeout value that the master router waits for all requests for a circuit before giving the permission for a router to take a circuit by issuing the `dls transparent timers` interface command. You can use the default values or you can create separate timeout values for NetBIOS and SNA sessions.

Figure 12-2 is a sample configuration of DLSw+ with the Ethernet redundancy feature. Routers A, B and C advertise their presence on the Ethernet via their Ethernet interfaces to the multicast MAC address 9999.9999.9999. Because Router B is the master, it keeps a database of all circuits handled within the domain and grants or denies permission for new circuit requests for Router A and Router C. There is no special configuration required for the end stations or for the remote peer. Only the DLSw+ devices on the LAN need the extra configuration. Master Router B waits 1.5 seconds after it receives the first IWANTIT primitive before assigning the new SNA circuit to one of the Ethernet redundancy peers because of the `dls transparent timers sna 1500` command.

Figure 12-2 Sample Diagram of DLSw+ with Ethernet Redundancy



```
Router A
dlsw local-peer peer id 10.2.24.2
dlsw remote-peer 0 tcp 10.2.17.1
int e1
ip address 150.150.2.1 255.255.255.0
dlsw transparent redundancy-enable 9999.9999.9999
```

```
Router B
dlsw local-peer peer-id 10.2.24.3
dlsw remote-peer 0 tcp 10.2.17.1
int e1
ip address 150.150.2.3 255.255.255.0
dlsw transparent redundancy-enable 9999.9999.9999 master-priority 1
dlsw transparent timers sna 1500
```

```
Router C
dlsw local-peer peer-id 10.2.24.4
dlsw remote-peer 0 tcp 10.2.17.1
int e1
ip address 150.150.2.4 255.255.255.0
dlsw transparent redundancy-enable 9999.9999.9999
```

```
Router D
dlsw local-peer peer-id 10.2.17.1 promiscuous
```

Verifying the Ethernet Redundancy Feature

There are several show commands that enable the user to verify the configuration for Ethernet redundancy. Verify that the master router is configured correctly by issuing the show dlsw transparent neighbor command on the appropriate routers. The following sample shows output from the show dlsw transparent neighbor command:

```
routerB#show dlsw transparent neighbor
Interface E0
0006.e278.6c0e SELF Master
0009.fa50.0b1c Rcvd Master-Accepted VALID
```

The output shows that Router B is the master router whose MAC address is 0006.e278.6c0e. The other router, with a MAC address of 0009.fa50.0b1c, is a slave router on the common domain. The master router received a packet from the slave and notes the router is VALID

Verify that the cache of the routers is populating correctly by issuing the show dlsw reachability command on the Ethernet redundancy routers. The following sample shows output from the show dlsw reachability command:

```
DLSw Local MAC address reachability cache list
Mac Addr      status      Loc.      port      rif
0004.f557.c156 FOUND      LOCAL    Ethernet1  --no rif--
DLSw Remote MAC address reachability cache list
Mac Addr      status      Loc.      peer
0004.f557.c164 FOUND      REMOTE   10.2.17.1(2065) max-lf(1500)
DLSw Local NetBIOS Name reachability cache list
NetBIOS Name  status      Loc.      port      rif
stationx      FOUND      LOCAL    Ethernet1  --no rif--
DLSw Remote NetBIOS Name reachability cache list
NetBIOS Name  status      Loc.      peer
stationy      FOUND      REMOTE   10.2.17.1(2065) max-lf(17800)
```

The output shows that End Station X (0004.f557.c156) is reachable directly through interface Ethernet 1. Verify that the master router has the correct circuits in its cache by issuing the show dlsw transparent cache command. The following sample shows output from the show dlsw transparent cache command issued on Master Router B:

```
routerB#show dlsw transparent cache
Interface Ethernet0/1
Circuit Cache
local addr(lsap)  remote addr(dsap)  state      Owner
0000.3028.92b6(08) 0007.0db1.238c(08) POSITIVE   SELF
0000.3028.92b6(08) 0008.dec3.609e(12) NEGATIVE   0009.fa50.0b1c
Total number of circuits in the Cache:2
```

The output shows that there are two circuits in the transparent bridging domain. The first circuit listed is "owned" by Router B, and all traffic over the circuit flows through Router B. The second circuit listed has been granted to the redundant peer with the MAC address 0009.fa50.0b1c, and all traffic over that circuit flows through the corresponding router. The "state" column indicates to which peer the circuit belongs.



Verify the number of circuits being handled by each of those peers by issuing the show dlsw peer command on the Ethernet redundancy routers. The following sample shows output from the show dlsw peer command:

```
router A#show dlsw peer
Peers:                state      pkts_rx  pkts_tx  type  drops  ckts
TCP  uptime
  TCP 10.2.17.1        CONNECT    4936     61068   conf    0     0     0
16:17:13
Total number of connected peers:1
Total number of connections:  1

router B#show dlsw peer
Peers:                state      pkts_rx  pkts_tx  type  drops  ckts
TCP  uptime
  TCP 10.2.17.1        CONNECT   748975  15817022 conf    0     1     0
16:15:47
Total number of connected peers:1
Total number of connections:  1

router C#show dlsw peer
Peers:                state      pkts_rx  pkts_tx  type  drops  ckts
TCP  uptime
  TCP 10.2.17.1        CONNECT    7387     104617  conf    0     1     5
16:17:04
Total number of connected peers:1
Total number of connections:  1

router D#show dlsw peer
Peers:                state      pkts_rx  pkts_tx  type  drops  ckts
TCP  uptime
  TCP 10.2.24.2        CONNECT    61068     4936   prom    0     0     0
16:17:14
  TCP 10.2.24.3        CONNECT   15817022  748975  prom    0     1     0
16:15:47
  TCP 10.2.24.4        CONNECT   104617     7387   prom    0     1     0
16:17:04
Total number of connected peers:3
Total number of connections:  3
```

The output shows that redundant peers A, B, and C all are connected to remote Router D. Currently there are two circuits, one between Router B and Router D, and the other between Router C and Router D.

Configuration Considerations with Ethernet Redundancy

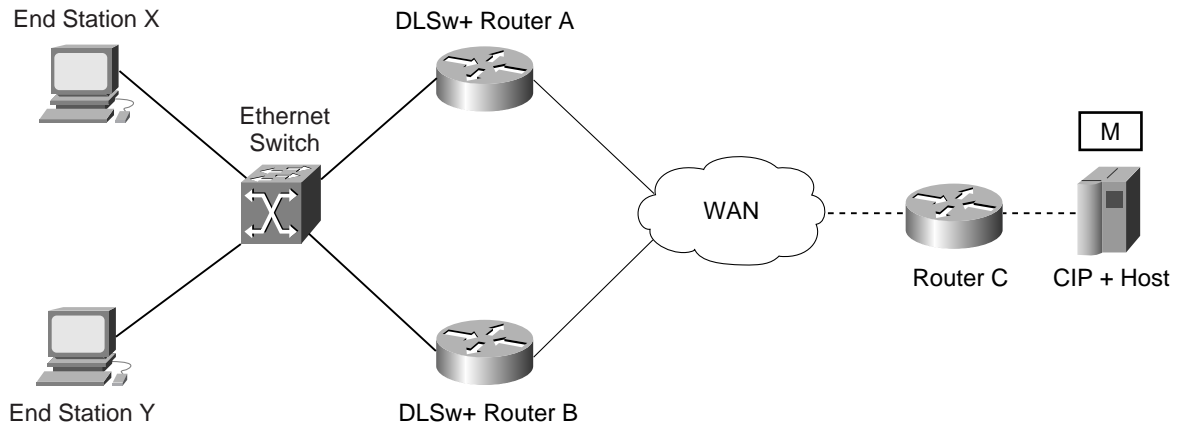
Because of issues with the propagation of UI frames, NetBIOS browsing is not supported in this release. Users must know the NetBIOS name of the server to which they wish to connect.

Do not configure both the global dlsw bridge-group command (in order to bridge to another LAN interface) and the interface dlsw transparent-redundancy command on the same device. You can, however, configure the interface bridge-group command on non-Ethernet redundancy interfaces. This type of configuration means that no bridging occurs between the two groups. Also, do not configure transparent bridging on an Ethernet interface that is configured for Ethernet redundancy. You can, however, configure two separate interfaces (on separate Ethernet segments) on the same router being able to reach the same MAC address with either the Ethernet redundancy feature or by using the dlsw bridge-group command (using separate bridge groups).

Support Switch Devices in an Ethernet Environment

Ethernet redundancy with switch devices requires further changes because of the way in which switches handle and direct traffic. Switches direct traffic by observing a frame's SMAC and by observing from which port the frame arrives. They forward all traffic to a particular address from a specific port rather than flooding all of its ports. In a normal Ethernet environment, this method is sufficient because there can only be one unique path to any MAC address. However, this method does not work in an environment where there are multiple DLSw+ routers on the same transparent bridged domain that can reach the same MAC address (see Figure 12-3).

Figure 12-3 DLSw+ with Ethernet Redundancy in a Switched Environment




Because Routers A and B are hooked to different ports on the Ethernet switch, the switch sees traffic from one SMAC coming into multiple ports. The Ethernet switch thinks the MAC address of the Host appears on two different places on a LAN. This design breaks the Ethernet rule of having only one path to any MAC address. It gives the appearance of a bridging loop that the Spanning-Tree Protocol did not resolve. Because SNA is connection oriented, the session is eventually torn down.

The DLSw+ Ethernet redundancy feature solves this problem with MAC address mapping. MAC address mapping ensures that a particular SMAC is seen by the switch on only one port at a time. Furthermore, the routers monitor each other's MAC address mapping so that they adequately serve as each other's back up in the case of a router failure.

In Figure 12-4, Router A is configured to map MAC address M' (M prime) to MAC address M and Router B is configured to map MAC address M'' (M double prime) to MAC address M. End Station X is configured to use M' as its SNA DMAC and End Station Y is configured to use M'' as its SNA DMAC.

Figure 12-4 DLSw+ Network Showing MAC Address Mapping



The following sequence occurs:

- End Station X sends out a TEST poll searching for its DMAC M'.
- The switch floods the request to all its ports because it is a new circuit and the switch has not heard of M'.
- The switch notes the port through which End Station X can be reached.
- Router A sees the TEST poll and recognizes that it is mapping M' to MAC address M. It replies to the switch with a TEST final.
- The switch populates its cache with M' and, because it knows where End Station X is now, the frame is directed out a single port rather than flooded to all its ports.
- The end station sends an XID because it is ready to start an SNA session.
- The Ethernet switch directs the frame out the port to which Router A is attached because it has seen a packet with the SMAC of M'.
- Router A asks Router B permission to take the circuit since Router B is the master.
- Router A receives permission to take the circuit because it is doing MAC address mapping for M' to M. Router A sends a CANUREACH_cs to begin the process of creating a circuit.
- Router A does MAC address translation by replacing M' in the DMAC field with M, the actual MAC address of the mainframe resource. From this point forward, any frames directed from Router A toward the WAN are referred to as M and any frames being directed from Router A toward the LAN are referred to as M'. Some level of load balancing is achieved if half of the end stations are configured to use DMAC M' and the other half are configured to use M''.

In the case of a router failure, the other router detects the failure and seamlessly takes over the failed router's mapping responsibilities. In Figure 12-4, if Router A fails, the switch thinks it can still reach MAC address M' out the port that is connected to failed Router A. Router B takes over the mapping responsibilities for Router A by sending a TEST frame with SMAC M' and a multicast DMAC to the switch. The switch notes the SMAC M' and assumes the resource moved and updates its CAM table appropriately. Now End Station X tries to reestablish its connection to the mainframe by sending out an XID poll to its configured DMAC M'. The switch knows to direct this frame out the port to which Router B is attached because of the TEST frame Router B sent earlier. Router B assumes the mapping responsibilities of Router A by mapping M' to M and continues its own mapping responsibilities of mapping M'' to M.

When Router A recovers, Master Router B realizes that Router A should be mapping M' to M. Both routers cannot map M' to M simultaneously because the switch cannot handle multiple ports with reachability to the same MAC address. Master Router B, therefore, stops mapping M' to M and the existing sessions are taken down and recovered through Router A.


Configuring Ethernet Redundancy in a Switched Environment

To configure Ethernet redundancy in a switched environment, you need to:

- Step 1. Configure the `dlsw transparent redundancy-enable` interface command on those DLSw+ devices located on the same transparent bridged domain. Configure the multicast MAC address to which all the devices will advertise their presence. It is at this point that you would also elect the master router.
- Step 2. (Optional) Configure the `dlsw transparent timers` interface command to set the amount of time that the master router waits for all requests for a circuit before giving permission to a router to accept a circuit. You can use the default values or you can create separate timeout values for NetBIOS and SNA sessions.
- Step 3. Issue the `dlsw transparent switch-support` global configuration command on the routers connected to the Ethernet switch.
- Step 4. Issue the `dlsw transparent map` interface command on the Ethernet interfaces that are connected to the Ethernet switch to enable MAC address mapping. The user creates the local mac address which maps to the actual MAC address that is specified in the remote mac option. It is recommended that you configure a backup to the DLSw+ device that will resume the mapping responsibilities if the DLSw+ device fails.

Figure 12-5 is a sample configuration of the DLSw+ Ethernet redundancy feature in a switched environment. The Ethernet switch sees the device with MAC address 4000.0010.0001 one port at a time because Router A and Router B have mapped different MAC addresses to it. This configuration is known as MAC address mapping. Router A is configured so that noncanonical MAC address 4000.0001.0000 (canonical MAC address 0200.0080.0000 as configured in the router) maps to the actual device with noncanonical MAC address 4000.0010.0001 (canonical MAC address 0200.0008.0080). Router B is configured so that noncanonical MAC address 4000.0201.0001 (canonical MAC address 0200.4080.0080) maps to the actual device with noncanonical MAC address 4000.0010.0001 (canonical MAC address 0200.0008.0080). Router A and B backup one another. Router A is configured as the master with a default priority of 100. Master Router B waits 1.5 seconds after it receives the first IWANTIT primitive before assigning the new SNA circuit to one of the Ethernet redundancy peers because of the `dlsw transparent timers sna 1500` command.

Figure 12-5 DLSw+ with Ethernet Redundancy in a Switched Environment



Router A

```
dlsw local peer peer-id 10.2.17.1
dlsw remote-peer 0 tcp 10.3.2.1
dlsw transparent switch-support
int e 0
mac-address 4000.0000.0001
ip address 150.150.2.1 255.255.255.0
dlsw transparent redundancy-enable 9999.9999.9999 master-priority 1
dlsw transparent map local-mac 0200.0080.0000 remote-mac 0200.0008.0080 neighbor 4000.0000.0011
dlsw transparent timers sna 1500
```

Router B

```
dlsw local peer peer-id 10.2.17.2 promiscuous
dlsw transport switch-support
int e 1
mac-address 4000.0000.0011
ip address 150.150.3.1 255.255.255.0
dlsw transparent redundancy-enable 9999.9999.9999
dlsw transparent local-mac 0200.4080.0080 remote-mac 0200.0008.0080 neighbor 4000.0000.0001
```

Verifying the Ethernet Redundancy Feature in a Switched Environment

Use the `show dlsw transparent map` command (in addition to the other commands listed in the “Verifying the Ethernet Redundancy feature” section of this chapter) to verify the configuration in a switched environment.

1. Verify that the created MAC address to which all the Ethernet redundancy routers are mapped is configured correctly by issuing the `show dlsw transparent map` command on all the routers configured for Ethernet redundancy. The command should be issued on all the routers configured for the Ethernet redundancy feature to ensure that the local MAC addresses match. By viewing the output you can also verify that a router is configured to backup another router’s MAC address mapping functions.

The output from Router A and Router B shows the created MAC addresses are 4000.0001.0000 and 4000.0201.0001:

```
router A#show dlsw transparent map
Interface Vlan200
  LOCAL Mac          REMOTE MAC          BACKUP
  -----
  4000.0001.0000     4000.0010.0001     0200.0000.0088     STATIC
  4000.0201.0001     4000.0010.0001     0200.0000.0088     DYNAMIC(Passive)
```

```
router B#show dlsw transparent map
Interface Vlan200
  LOCAL Mac          REMOTE MAC          BACKUP
  -----
  4000.0201.0001     4000.0010.0001     0200.0000.0080     STATIC
  4000.0001.0000     4000.0010.0001     0200.0000.0080     DYNAMIC(Passive)
```

Configuration Considerations with Ethernet Redundancy in a Switched Environment

DLSw+ local switching is not supported between two Ethernet redundancy interfaces, or between an Ethernet redundancy interface and any other LAN-type media (Token Ring, ISL, LANE, or FDDI).