

Getting Started

This chapter describes the basic configuration commands required for a DLSw+ network. It begins with a description of the minimum required configuration and then provides examples for Token Ring, Ethernet, SDLC, and QLLC environments. This section assumes that you are familiar with basic router configuration.

Minimum Required Configuration


Configuring DLSw+ on most networks is not difficult. Every router that supports DLSw+ must have a **dlsw local-peer** command; **dlsw remote-peer** commands are optional, but usually at least one side of a peer connection must configure a remote peer. If a DLSw+ peer configuration omits **dlsw remote-peer** commands, the **dlsw local-peer** command must specify the **promiscuous** keyword. Promiscuous routers will accept peer connection requests from routers that are not preconfigured. This feature allows you to minimize changes to central site routers when branch offices are added or deleted. It also minimizes required coordination of configurations.

If you have used RSRB in the past, you need to know what *not* to configure. With DLSw+, you do not need proxy explorer, NetBIOS name caching, SDLC-to-LLC2 conversion (SDLLC), or source-route translational bridging (SR/TLB). All of these features are built into DLSw+.

Note: SR/TLB is required, however, when doing local translation between Ethernet and Token Ring.

In Figure 2-1, the branch router specifies both a **dlsw local-peer** and a **dlsw remote-peer** command. The headquarters router specifies only a **dlsw local-peer** command, but it specifies **promiscuous** on the **dlsw local-peer** command to allow it to dynamically accept connections from branch routers. The peer ID specified on the **dlsw local-peer** command is the loopback address configured via the **interface loopback 0** command or the IP address associated with a specific LAN or WAN interface. However, if you use a LAN or WAN IP address, the interface must be up for DLSw+ to work.

Figure 2-1 Example of **dls w local-peer** and **dls w remote-peer** Commands



The number following **dls w remote-peer** is the ring list number. Ring lists are an advanced topic, so for now, specify zero in this space, which indicates that ring lists are not in use. There are other options on the **dls w local-peer** and **dls w remote-peer** commands, but they are not required. These options are covered in the “Advanced Features” chapter.

In addition to specifying local and remote peers, you must map the following local data-link controls to DLSw+:

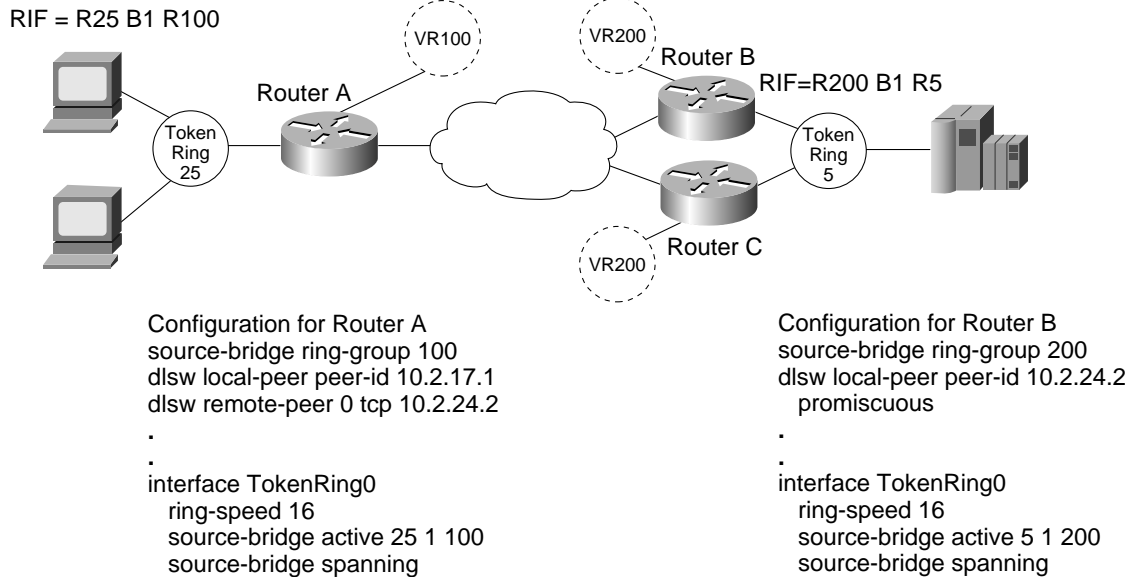
- *Token Ring*—Define a virtual ring using the global **source-bridge ring-group** command and include a **source-bridge** command that tells the router to bridge from the external Token Ring to that virtual ring on the interface
- *Ethernet*—Map a specific Ethernet bridge group to DLSw+
- *SDLC*—Define the SDLC devices and map the SDLC addresses to DLSw+ virtual MAC addresses
- *QLLC*—Define the X.25 devices and map the X.25 addresses to DLSw+ virtual MAC addresses
- *FDDI*—Define a virtual ring using the global **source-bridge ring-group** command and include an SRB statement that tells the router to bridge from the external FDDI to that virtual ring; FDDI is supported via SRB in Cisco IOS Release 11.2

The rest of this chapter provides sample configurations for Token Ring, Ethernet, SDLC, and QLLC.

Token Ring

Figure 2-2 shows a sample DLSw+ configuration for Token Ring. Traffic that originates on Token Ring is source-route bridged from the local ring onto a source-bridge ring group and then picked up by DLSw+. You must include a global **source-bridge ring-group** command that specifies a virtual ring number. In addition, you must include an interface **source-bridge** command that tells the router to bridge from the physical Token Ring to the virtual ring.

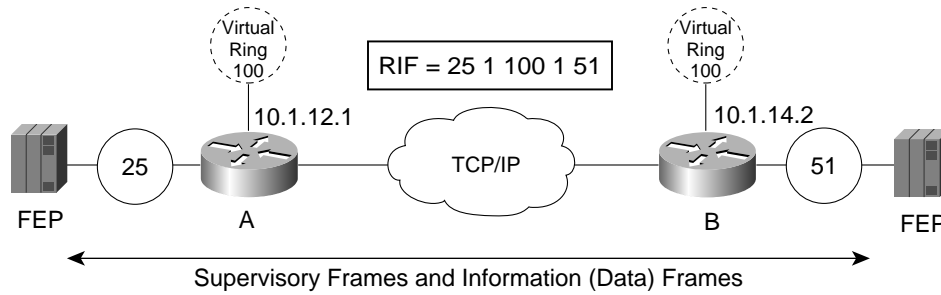
Figure 2-2 Simple Token Ring DLSw+ Configuration



DLSw+ supports RIF termination, which means that all remote devices appear to be attached to the virtual ring specified in the interface **source-bridge** command. In Figure 2-2 from the host end, all the devices attached to Router A would appear to reside on Virtual Ring 200. Conversely, from the remote site, the FEP would appear to reside on Virtual Ring 100. As illustrated in this figure, the virtual rings specified in peer routers do not have to match. If multiple routers are attached to the same physical ring, as shown in Routers B and C, by specifying the same ring group number in each of them, you can prevent explorers from coming in from the WAN and being forwarded back onto the WAN.

DLSw+ also supports TCP/IP with RIF Passthru, which means the RIF is not terminated and the entire source-route bridged path appears in the RIF. This solution is used to allow multiple active paths between FEP's. For proper configuration of the TCP/IP with RIF Passthru, the virtual ring numbers must match between the DLSw+ peers, however, the Token Ring numbers must be unique (see Figure 2-3).

Figure 2-3 DLSw+ Network Configured with TCP/IP RIF Passthru



```

Router A
source-bridge ring-group 100
dlsw local-peer peer-id 10.1.12.1
dlsw remote-peer 0 tcp 10.1.14.2 rif-passthru 100
interface tokenring 0
  ring-speed 16
  source-bridge 25 1 100
  
```

```

Router B
source-bridge ring-group 100
dlsw local-peer peer-id 10.1.14.2
dlsw remote-peer 0 tcp 10.1.12.1 rif-passthru 100
interface tokenring 0
  ring-speed 16
  source-bridge 51 1 100
  source-bridge spanning
  
```

Ethernet


Traffic that originates on Ethernet is picked up from the local Ethernet bridge group and transported across the DLSw+ network. DLSw+ always transfers data in noncanonical format. In Figure 2-4, you do not need to configure the left router for translational bridging or worry about what media resides on the other side of the WAN. DLSw+ will automatically make the correct MAC address conversion depending on the destination media. When DLSw+ receives a MAC address from an Ethernet-attached device, it assumes it is canonical and converts it to noncanonical for transport to the remote peer. At the remote peer, the address is either passed unchanged to Token Ring-attached end systems or converted back to canonical if the destination media is Ethernet. Note that when an SNA resource resides on Ethernet, if you configure a destination SNA address in that device, you must use canonical format. For example, Ethernet-attached IBM 3174s must specify the MAC address of the FEP in canonical format. If the Token Ring or noncanonical format of the MAC address of the FEP is 4000.3745.0001, the canonical format is 0200.ECA2.0080.

Note: Some environments avoid this issue by using MAC addresses consisting of only “magic numbers”—numbers that are the same in canonical and noncanonical formats. These numbers are 00, 18, 24, 3C, 42, 5A, 66, 7E, 81, 99, A5, BD, C3, DB, E7, and FF.



In Figure 2-4, the data is transferred directly to a Cisco router with a CIP, but it could be any DLSw-compliant router, and the upstream SNA end system could reside on any supported media.

Figure 2-4 Simple Ethernet DLSw+ Configuration



```
dlsw local-peer peer-id 10.2.17.1
dlsw remote-peer 0 tcp 10.2.24.2
.
.
dlsw bridge-group 1
interface Ethernet0
  no ip address
  bridge-group 1
  bridge 1 protocol ieee
```

```
source-bridge ring-group 200
dlsw local-peer peer-id 10.2.24.2
promiscuous
.
.
interface channel 0/1
  csna 0100 40
  csna 0100 41
  int chan 0/2
  lan tokenring 0
  source-bridge 1000 1 200
  adapter 0 4000.0000.0401
  adapter 1 4000.0000.0403
```


SDLC

Configuring SDLC devices is a bit more complicated. For SDLC devices, you must know whether the device is a PU 1, PU 2.0, or PU 2.1. For PU 2.0 devices, you must know the IDBLK and IDNUM that was specified in the Virtual Telecommunications Access Method (VTAM) for that device, because the router plays a greater role in XID processing when SDLC PU 2.0 is involved. You must know if the router is the primary or secondary end of the SDLC line. In addition, if the attachment to the upstream SNA device is over a LAN, you must configure the MAC address of the destination upstream SNA device. In all cases, you must configure a virtual MAC address that will be mapped to an SDLC polling address.

Note: With SDLC scenarios, DLSw+ is recommended over serial tunnel (STUN) except when configuring PU 4-to-PU 4 SDLC-attached devices.

In Figure 2-5, the SDLC-attached devices are each given a common base virtual MAC address of 4000.3174.0000. The router will replace the last two digits of the virtual MAC address with the SDLC address of the device. The device at SDLC address C1 appears to have MAC address 4000.3174.00C1, and the device at SDLC address C2 appears to have MAC address 4000.3174.00C2. In this example, both devices are PU 2.0 devices, so their XID must be configured and it must match what is specified as the IDBLK and IDNUM in VTAM. In addition, the router always assumes the primary role when attaching upstream from PU 2.0 devices.

Figure 2-5 Simple SDLC DLSw+ Configuration



```

Configuration for Router A
dlsw local-peer peer-id 10.2.17.1
dlsw remote-peer 0 tcp 10.2.24.2
interface serial 0
  encapsulation sdhc
  sdhc role primary
  sdhc vmac 4000.3174.0000
  sdhc address c1
  sdhc xid c1 01712345
  sdhc partner 4000.3745.0001 c1
  sdhc dlsw c1

interface serial 1
  encapsulation sdhc
  sdhc role primary
  sdhc vmac 4000.3174.1000
  sdhc address c2
  sdhc xid c2 01767890
  sdhc partner 4000.3745.0001 c2
  sdhc dlsw c2
  
```


The router can be the secondary end of an SDLC line (for example, when connecting to a FEP over SDLC). In this case, specify **secondary** in the interface **sdhc role** command, and for PU 2.1 devices, specify **xid-passthru** in the interface **sdhc address** command.

In Cisco IOS Release 11.0 and later, DLSw+ supports multidrop PU 2.0/2.1. In Figure 2-6, the multidrop PU 2.0 configuration includes an interface **sdhc xid** command for each PU 2.0 device.

For multidrop lines with a mix of PU 2.1 and 2.0 devices, specify **primary** in the interface **sdhc role** command. For PU 2.0 devices, you must code the IDBLK and IDNUM in the interface **sdhc xid** command. For PU 2.1 devices, you can omit the **sdhc xid** command. However, in the **sdhc address** command, you need to specify **xid-poll**.

Alternately, when all devices on a line are PU 2.1, you can specify **sdhc role prim-xid-poll**, in which case you do not need to specify **xid-poll** in each **sdhc address** command.

Figure 2-6 Multidrop SDLC DLSw+ Configuration



Configuration for Router A
Both C1 and C2 are PU 2.0

```
dlsw local-peer peer-id 10.2.17.1
dlsw remote-peer 0 tcp 10.2.24.2
interface serial 0
mtu 4400
no ip address
encapsulation sdhc
no keepalive
clockrate 19200
sdhc role primary
sdhc vmac 4000.3174.0000
sdhc address c1
sdhc xid c1 01712345
sdhc partner 4000.3745.0001 c1
sdhc address c2
sdhc xid c2 01767890
sdhc partner 4000.3745.0001 c2
sdhc dlsw c1 c2
```

Configuration for Router A, Mixed PU
2.0 and 2.1


```
interface serial 0
...
sdhc role primary
sdhc vmac 4000.3174.0000
sdhc address c1 xid-poll
sdhc partner 4000.3745.0001 c1
sdhc address c2
sdhc xid c2 01767890
sdhc partner 4000.3745.0001 c2
sdhc dlsw c1 c2
```

Configuration for Router A, All PU 2.1

```
interface serial 0
...
sdhc role prim-xid-poll
sdhc vmac 4000.3174.0000
sdhc address c1
sdhc partner 4000.3745.0001 c1
sdhc address c2
sdhc partner 4000.3745.0001 c2
sdhc dlsw c1 c2
```

In Cisco IOS Release 11.3 and later, DLSw+ supports LLC2 to SDLC between PU 4 devices (see Figure 2-7).

Figure 2-7 LLC2-to-SDLC DLSw+ Configuration



```
Router A
interface serial 0
  mtu 4096
  ip address 10.4.21.2 255.255.255.0
  encapsulation frame-relay IETF
  keepalive 12
  frame-relay map 11c2 46
  frame-relay map 11c2 45
  frame-relay map ip 10.4.21.1 43 broadcast
  frame-relay map ip 10.4.21.3 45 broadcast
  frame-relay map ip 10.4.21.4 46 broadcast
  frame-relay lmi-type ansi

interface tokenring 0
  mac-address 4000.1250.1001
  no ip address
  ring-speed 16
  fras map 11c 4000.1060.1000 4 4 Serial0 frame-relay 45 4 4
```

```
Router B
interface serial 0
  mtu 4096
  ip address 10.4.21.3 255.255.255.0
  encapsulation frame-relay IETF
  keepalive 12
  no fair-queue
  frame-relay map 11c2 53
  frame-relay map 11c2 54
  frame-relay map 11c2 56
  frame-relay map ip 10.4.21.1 53 broadcast
  frame-relay map ip 10.4.21.3 54 broadcast
  frame-relay map ip 10.4.21.4 56 broadcast
  frame-relay lmi-type ansi

interface serial 1
  no ip address
  encapsulation sdhc
  no keepalive
  clockrace 9600
  sdhc address 01 echo
  fras map sdhc 1 Serial0 frame-relay 54 4 4 fid4
```

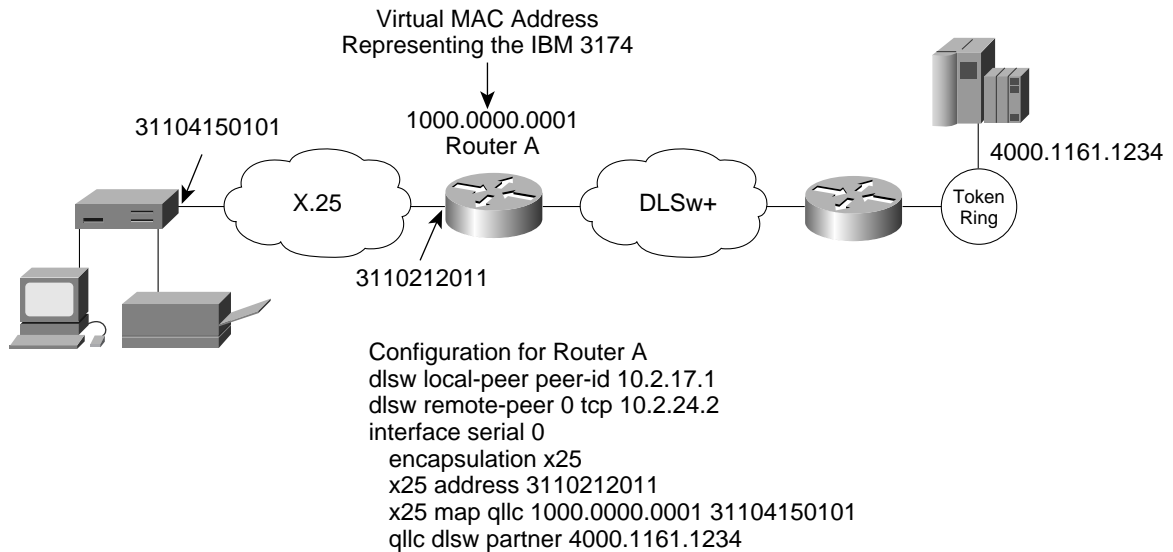
QLLC

QLLC is the data link used by SNA devices when connecting to X.25 networks. QLLC is a legacy protocol developed by IBM to allow the Network Control Program (NCP) to support remote connections over X.25. The software feature on NCP that supports QLLC is called Network Packet Switching Interface (NPSI). The QLLC protocol derives its name from using the Q-bit in the X.25 header to identify QLLC protocol primitives. QLLC essentially emulates SDLC over X.25.

Thus, DLSw+ performs QLLC conversion in a manner similar to SDLC conversion. Cisco's DLSw+ implementation added support for QLLC in Cisco IOS Release 11.0. Because QLLC is more complicated than Token Ring, Ethernet, or SDLC, three examples are included here.

Figure 2-8 shows DLSw+ being used to allow remote devices to connect to a DLSw+ network over an X.25 public packet switched network. In this example, all QLLC traffic is addressed to destination address 4000.1161.1234, which is the MAC address of the FEP. The remote X.25-attached IBM 3174 is given a virtual MAC address of 1000.0000.0001. This virtual MAC address is mapped to the X.121 address of the IBM 3174 (31104150101) in the X.25-attached router.

Figure 2-8 QLLC DLSw+ Configuration to a Single LAN-Attached Upstream Device

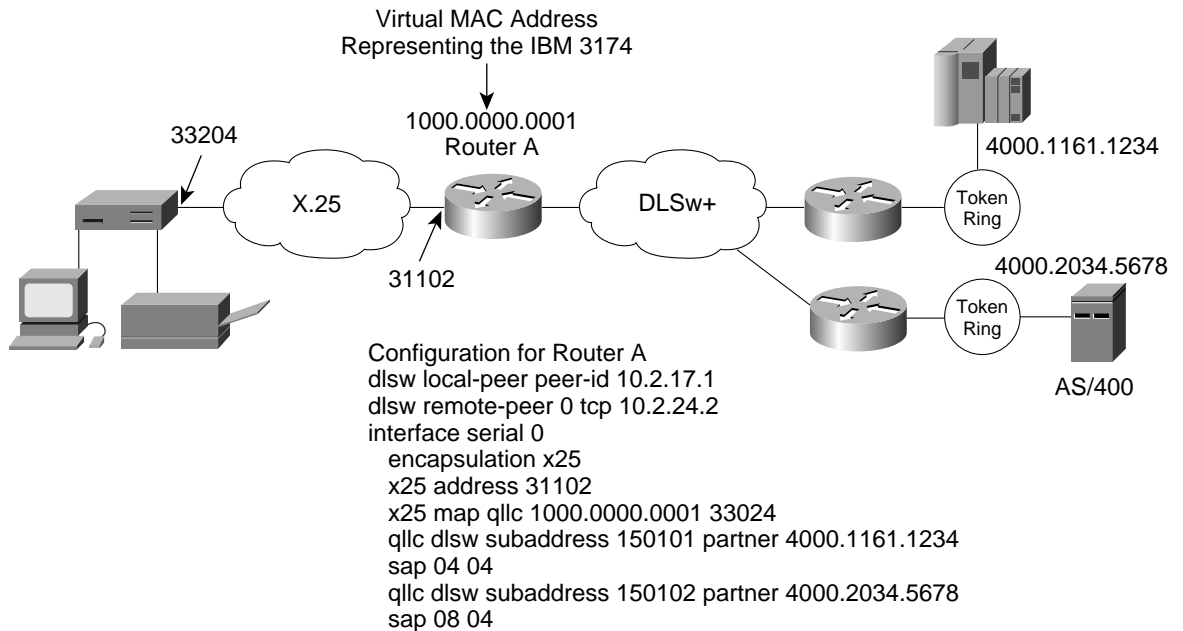


In Figure 2-9, a single IBM 3174 needs to communicate with both an AS/400 and a FEP. The FEP is associated with subaddress 150101, and the AS/400 is associated with subaddress 150102.

If an X.25 call comes in for 33204150101, the call is mapped to the FEP and forwarded to MAC address 4000.1161.1234. The IBM 3174 appears to the FEP as a Token Ring-attached resource with MAC address 1000.0000.0001. The IBM 3174 uses a source SAP of 04 when communicating with the FEP.

If an X.25 call comes in for 33204150102, the call is mapped to the AS/400 and forwarded to MAC address 4000.2034.5678. The IBM 3174 appears to the AS/400 as a Token Ring-attached resource with MAC address 1000.0000.0001. The IBM 3174 uses a source SAP of 08 when communicating with the AS/400.

Figure 2-9 QLLC DLSw+ Configuration for Support of Multiple Upstream LAN-Attached Devices



In Figure 2-10, two X.25 resources want to communicate over X.25 to the same FEP. In the router attached to the X.25 network, every X.25 connection request for X.121 address 31102150101 is directed to DLSw+. The interface **qlc dsw** command creates a pool of two virtual MAC addresses, starting with 1000.0000.0001. The first switched virtual circuit (SVC) established will be mapped to virtual MAC address 1000.0000.0001. The second SVC will be mapped to virtual MAC address 1000.0000.0002.

Figure 2-10 QLLC DLSw+ Configuration Supporting Multiple Downstream X.25-Attached Devices Communicating through an Upstream DLSw+ Network

