

RSRB Migration and Multivendor Interoperability

This chapter describes the differences between RSRB and DLSw+, as well as the following issues:

- The reasons you would migrate from RSRB to DLSw+
- The migration implications in terms of management, memory, and performance
- The steps to migrate from RSRB to DLSw+

In addition, this chapter describes interoperability with other RFC 1795 and RFC 2166 implementations, including valid configuration options.

RSRB and DLSw+ Comparison

RSRB, created in 1991, preceded DLSw+ and existed before there were routing standards. RSRB was the original Cisco implementation for transporting LLC2 traffic over an IP network. RSRB addressed a critical need in the market place, thus thousands of RSRB networks were built. DLSw+ has replaced many of these networks, but there are hundreds of RSRB networks still in existence, some with more than 1000 RSRB routers.

The AIW approved the first standard for SNA over IP in 1995. This standard was created in and was later documented in RFC 1795 and RFC 2166. DLSw+ complies with RFC 1795 and RFC 2166 and provides enhancements that allow DLSw+ networks to scale better and provide better availability than either RSRB or standard-only implementations. Most integrated SNA and IP networks that were installed since 1995 have been built using DLSw+.

DLSw+ includes functions that were previously provided in several other Cisco features, including RSRB, SDLC-to-LLC2 conversion (SDLLC), SR/TLB, proxy explorer, and NetBIOS name caching. Most environments using DLSw+ no longer need to configure any of these features.

Why Move to DLSw+?

Cisco has chosen DLSw+ as the strategic solution for SNA transport going forward. RSRB has not been enhanced since 1995 and no enhancements are planned. In addition, at some point in the future the new Cisco IOS Software releases will no longer include RSRB.

DLSw+ provides better functionality, manageability, and control than RSRB. DLSw+ addresses several RSRB limitations by including key functions such as local acknowledgment for devices on Ethernet and SDLLC for PU 2.1 devices. In addition, DLSw+ scales better than RSRB, is easier to configure and manage, and provides higher availability with load balancing and backup features. DLSw+ also offers multivendor interoperability. Table 8-1 illustrates the differences between RSRB and DLSw+.

Table 8-1 Comparison of Cisco RSRB to DLSw+

Benefits	RSRB Features	DLSw+ Features
Performance	IP load sharing Custom and priority queuing	IP load sharing Custom and priority queuing Circuit-level flow control ¹ Peer and port load sharing ¹
Availability	Nondisruptive rerouting around link failures Local acknowledgment on Token Ring and SDLC	Nondisruptive rerouting around link failures Local acknowledgment on Token Ring and SDLC Local acknowledgment on Ethernet ¹ Backup peers ¹ Fault tolerant and priority peers ¹
Scalability	Limited broadcast reduction	Broadcast reduction Dynamic peers ¹ UDP for UI frames ¹ RIF termination ¹ Broadcast optimization with peer groups and border peer caching ¹
Flexibility	Media conversion via SDLLC and SR/TLB (PU 2.0 only) SRB dynamics RIF Passthrough Transport options (FST, Direct) Support for end systems on Token Ring, SDLC (with SDLLC), or Ethernet (with SR/TLB) AST ² FST between unlike media via SDLLC ² LNM over FST ²	Media conversion built in (PU 2.0, 2.1 and PU 4) SRB dynamics RIF termination or optional RIF passthrough Transport options (FST, direct) DLSw Lite (LLC2 encapsulation) Support for end systems on Token Ring LANE, Token Ring ISL, and SRB FDDI Capabilities exchange Peer biasing with cost SNA DDR Promiscuous peers Multivendor interoperability

1. Supported by DLSw+ but not by RSRB

2. Supported by RSRB but not by DLSw+

Cisco designed DLSw+ in a modular fashion to maximize stability and to facilitate new feature additions. The circuit concept in DLSw simplifies management. The Cisco implementation protects your investment in the technology and simplifies network integration of acquired companies because it can interoperate with other standard-compliant implementations. Finally, DLSw+ surpasses RSRB as the most commonly employed technique for SNA and client/server integration.

Possible Migration Inhibitors

A few environments will not be able to move from RSRB to DLSw+ at this time. They may be on older software releases and require features that were added to DLSw+ in a recent release of Cisco IOS Software. For example, RIF passthru was added in Cisco IOS Release 12.0 and is required for FEP-to-FEP communication over parallel SRB paths.

There are a few RSRB features not available in DLSw+ or planned for future releases, including the following:

- FST between SDLC and LANs
- LAN Network Manager over FST
- Automatic Spanning Tree (AST), which is used by source-route bridges to determine whether they should forward single-route explorers



Migration Considerations

The first two questions people ask when considering migration are:

- Does DLSw+ perform as well as RSRB?
- Does DLSw+ require additional memory?

From a performance standpoint, DLSw+ uses the same or slightly fewer CPU cycles to handle an equivalent amount of traffic. (This comparison assumes that either local acknowledgment is turned on for both or off for both.) However, DLSw+ uses more memory than RSRB. The key reason DLSw+ requires more memory is that DLSw+ maintains state information for every circuit and caches entries for multiple active paths. Maintaining state information simplifies management, and maintaining cache entries allows better network design. Even with the additional memory requirements, most networks run well with the default memory that comes with the router and software subset. For example, the Cisco 2500 Series router (branch router) in a typical branch environment with 20 to 40 PUs and LUs and the standard memory configuration that comes with any of the IBM images, runs DLSw+ quite well. If you are running RSRB with an older level of the Cisco IOS Software, you may want to verify that your current routers can support DLSw+ with the memory they have. The Cisco IOS Software subset image takes up the bulk of the memory, and the image size has grown over time. The memory that is required to store the image is the most important part of the equation. (The size of any Cisco IOS Software feature set varies by release, so that information is not included here.) If necessary, you can approximate the memory required by DLSw+ from the formulas provided in Appendix A.

Migration Options

There are several ways to migrate to DLSw+. Which migration option you use depends on how the current RSRB peering structure is set up, whether your RSRB network allows any-to-any communication, and which design you want to use for your DLSw+ network. This section describes the following options:

- Migrating hierarchical networks (where all communication is from remote routers back to one or more central site routers) using
 - Separate routers for RSRB and DLSw+ peering
 - RSRB and DLSw+ concurrently in the same data center peer
- Migrating any-to-any networks from a
 - Fully meshed RSRB network (where every router is peered to every other router) to a fully meshed DLSw+ network
 - Fully meshed RSRB network to DLSw+ peer groups
 - Multihop RSRB network to DLSw+ peer groups


Before you can decide which of these options is best for your environment, it is essential to understand your current RSRB network.

Understanding Your Current RSRB Network Topology

To migrate an existing RSRB network to DLSw+, you must first understand what connectivity your RSRB network provides. Your RSRB network might enable communication that is not obvious from your network definitions.

For example, as shown in Figure 8-1, if Router A has RSRB connections to Router B and Router C, traffic might flow from Router B to Router C even though they are not peers. This situation occurs only if Router A is configured with two virtual rings and a separate physical interface connecting Physical Ring 10 with each virtual ring. Identify whether or not there are multiple ring-group numbers in use within the same part of an RSRB network.

Figure 8-1 Sample RSRB Network with Different Ring Numbers




To understand how this router configuration affects the data paths in the network, it is helpful to understand the use of the RIF field in explorer frames used for path discovery. When a device sends a path discovery frame on an SRB media, it usually inserts a RIF indicating that this frame is an explorer frame. When a device propagates an explorer frame through the network, source-route bridges on the network copy the explorer off the ring from which it originated onto one or more other rings that the bridge connects. As the bridge copies the frame, it modifies the RIF, indicating the path that the explorer frame took. The RIF lists every ring and every bridge that this frame has crossed. Hence, by looking at the RIF, it is possible to uniquely identify the entire path a frame has taken through the SRB network.

The bridge checks the existing RIF to ensure that the explorer has not already traversed its ring before it places a new copy of an explorer frame on its ring. This action prevents an explorer from looping around an SRB network until it exceeds the maximum number of permitted bridge “hops” (seven in most SRB implementations).

Many SRB designs make use of this feature by defining all of their virtual ring numbers as the same value. This action ensures that no frame traverses any RSRB hop more than once. In Figure 8-2, Router X peers to Router A and Router Y peers to Router B. If Router A forwards an explorer from Router B onto Ring 10, SRB in Router B does not pick up that explorer. The frame cannot be copied to the virtual ring in Router B because the RIF shows that the explorer has already traversed Ring 100. Using the same virtual ring number for RSRB peers in this manner helps limit the total number of explorers traversing a WAN.

Figure 8-2 RSRB Network with the Same Virtual Ring Numbers



However, if RSRB peers use different virtual ring numbers, an explorer (or, in fact, any frame) might traverse two different RSRB hops. For example, in Figure 8-1, the explorers originating in Router B go through Router A, onto the physical ring attached to Router A, and then back through Router A to Router C. The key point to remember is that data paths through the network are not obvious from a simple examination of the peer definitions. You need to look at ring-group numbers and virtual ring numbers.


To quickly determine if your network has this characteristic, look for an RSRB router using two different ring-group numbers, as shown in the following configuration:

```
source-bridge ring-group 100
source-bridge remote-peer 100 tcp 1.1.1.1
source-bridge remote-peer 100 tcp 1.1.2.1
source-bridge ring-group 200
source-bridge remote-peer 200 tcp 1.1.1.1
source-bridge remote-peer 200 tcp 1.1.3.1
!
interface TokenRing 0
ip address 1.1.1.1 255.0.0.0
ring-speed 16
source-bridge 10 1 100
source-bridge spanning
!
interface TokenRing 1
no ip address
ring-speed 16
source-bridge 10 1 200
source-bridge spanning
```

Looking again at Figure 8-1 and the preceding configuration, both interfaces TokenRing 0 and TokenRing 1 connect to the same physical ring (Ring 10). This configuration allows two different ring-groups to share data. Data received from peer 1.1.2.1 on ring-group 100 would be put onto ring 10 through interface TokenRing 0. It would then be transferred into ring-group 200 through interface TokenRing 1 and sent to peer 1.1.3.1. Because there is a LAN hop in between (Ring 10) and because two separate ring-groups were used, peers 1.1.2.1 and 1.1.3.1 are able to exchange data even though they do not peer to each other.

A similar situation occurs even if the two ring-groups do not exist in the same router. Assume a situation where Router A and Router B are both connected to the same physical Token Ring via interface TokenRing 0 as shown in Figure 8-3.

Figure 8-3 RSRB with Multiple Routers Connected to the Same Physical Ring



Assume the following configuration on Router A:


```
source-bridge ring-group 100
source-bridge remote-peer 100 tcp 1.1.1.1
source-bridge remote-peer 100 tcp 1.1.2.1
source-bridge remote-peer 100 tcp 1.1.3.1
!
interface TokenRing 0
ip address 1.1.1.1 255.255.255.0
ring-speed 16
source-bridge 10 1 100
source-bridge spanning
```

Assume the following configuration for Router B:

```
source-bridge ring-group 200
source-bridge remote-peer 200 tcp 1.1.1.2
source-bridge remote-peer 200 tcp 1.2.2.1
source-bridge remote-peer 200 tcp 1.2.3.1
!
interface TokenRing 0
ip address 1.1.1.2 255.255.255.0
ring-speed 16
source-bridge 10 1 200
source-bridge spanning
```

Traffic flows from any peer in ring-group 100 to any peer in ring-group 200 through physical Ring 10.

Both of these examples are common in RSRB configurations. Ring 10 in both examples is called an isolation (or de-encapsulation) ring. In both cases, the devices incur additional overhead because data travels through two RSRB hops to get from one remote branch router to another remote branch router. The data is encapsulated and de-encapsulated in TCP/IP twice. In addition, if local acknowledgment is used, then the LLC2 sessions must be terminated an extra time. Often, this configuration is selected to combine two existing RSRB networks because it is an easier solution than reconfiguring an entire RSRB network to match the ring-group number of another.



If either of these types of RSRB configurations exist in a network, then they are excellent candidates for DLSw+ border peers and peer groups. Using DLSw+ peer groups you can maintain the connectivity previously described without the overhead of multiple encapsulation steps. It provides the best of both worlds.

If neither of these conditions exists in your RSRB network, it should be possible to simply migrate by copying the existing RSRB peer definitions to equivalent DLSw+ peer definitions. Although this migration preserves existing data connectivity in the network, it might not result in the optimal network design.

Migrating Steps for a Hierarchical Network

In a hierarchical network, any-to-any connectivity is not required. Typically there are one or a small number of data center routers to which all remote sites must connect. These types of networks are the simplest to migrate. DLSw+ and RSRB can both run in the same router at the same time, but you may prefer to use a new router for the migration. This section describes both options.

Separate Routers for RSRB and DLSw+ Peering

One option for migrating a hierarchical RSRB network to DLSw+ is to put separate DLSw+ routers in parallel with the existing RSRB peering routers at each central site location. In certain situations, this is the only way to go (these situations are discussed later in this paper). It requires extra equipment, but only for the duration of the migration process, at which time the equipment can be redeployed.

To migrate a hierarchical RSRB network to a hierarchical DLSw+ network using separate, parallel routers, do the following:

- Step 1. Ensure that your routers are at Cisco IOS Release 10.3(2) or later (see Appendix B to determine which version of Cisco IOS Software you should install to get the features you want).
- Step 2. Configure a `dlsw local-peer` command, specifying the `promiscuous` keyword, at the new data center DLSw+ routers. If desired, you can replace the `promiscuous` keyword with `static dlsw remote-peer` command at the end of the migration process.
- Step 3. Select one remote site and delete the `source-bridge remote-peer` commands (and any related commands such as `SDLLC`, `SR/TLB`, `proxy explorer`, and `NetBIOS` name caching that are no longer required).
- Step 4. Add the appropriate `dlsw local-peer` command and one or more `dlsw remote-peer` commands (at that same remote site) that point to the central site DLSw+ routers.
- Step 5. Visit the RSRB device in the data center and remove the `source-bridge remote-peer` command that referred to the remote router you just modified.
- Step 6. Repeat Steps 3 through 5 with the remaining remote sites.

When all of the RSRB remote peers have been migrated to DLSw+, no remote peers will remain connected to the RSRB peering router. This router can now be removed and reused elsewhere in the network.

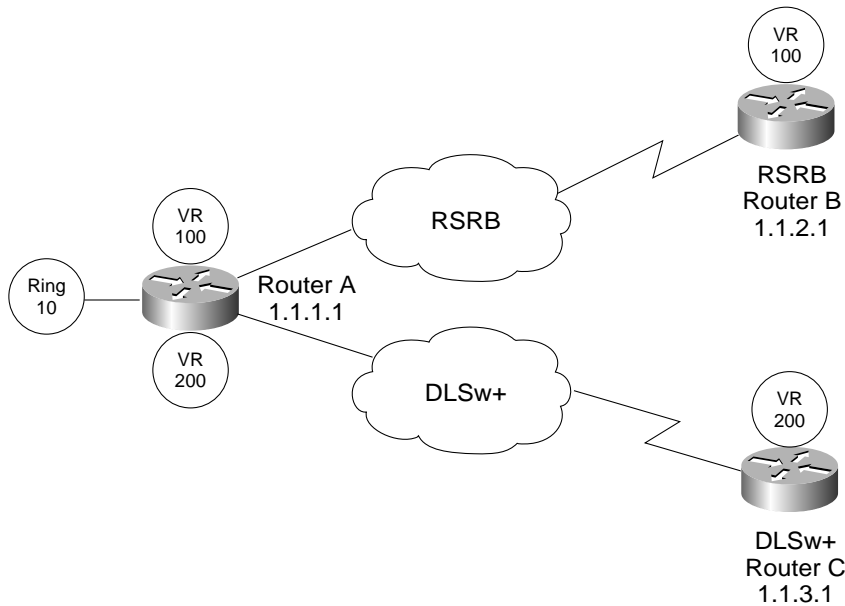
To keep explorers from going into the RSRB network from the DLSw+ network (or vice versa) during the course of the migration, the DLSw+ peers in the data centers should use the same `source-bridge ring-group` number as the RSRB peers with which they are being placed in parallel. The ring-group used at the remote sites does not matter as much, but it often is best to keep it the same, because this prevents explorer looping if there is an unknown back door data path. By reusing ring-group numbers, the SRB process recognizes packets that have already traversed the WAN via either RSRB or DLSw+ and discards these before they are passed back over the WAN.

RSRB and DLSw+ Concurrently in the Same Data Center Routers

Another option is to add DLSw+ to the existing RSRB peer devices running in the data center. DLSw+ was designed to run concurrently with RSRB in simple connectivity scenarios. However, technology limitations make this unfeasible in some situations. In general, between any pair of routers you should use either DLSw+ or RSRB, but not both, as shown in Figure 8-4. As in the previous example, migrate your RSRB network to DLSw+ one router at a time.

Figure 8-4 shows a sample central site configuration.

Figure 8-4 Central Site Router Configured to Communicate with Both an RSRB Router and a DLSw+ Router



Assume the following configuration on Router A:

```
source-bridge ring-group 100
source-bridge remote-peer 100 tcp 1.1.1.1
source-bridge remote-peer 100 tcp 1.1.2.1
!
dlsw local-peer peer-id 1.1.1.1
dlsw remote-peer 0 tcp 1.1.3.1
!
int tokenring 0
 source-bridge 10 1 100
 source-bridge spanning
```

To run DLSw+ and RSRB concurrently in data center peers, perform the following tasks:

- Step 1. Migrate your routers to Cisco IOS Release 10.3(2) or later (see Appendix B to determine which version of Cisco IOS Software you should install to get the features you want).
- Step 2. Configure a `dlsw local-peer` command using the same peering address as RSRB at the central site RSRB routers. For ease of migration, use the promiscuous keyword for the duration of the migration as discussed in the previous section.
- Step 3. Select one remote site and remove the devices `source-bridge remote-peer` commands (and any related commands such as `SDLLC`, `SR/TLB`, `proxy explorer`, and `NetBIOS` name caching that are no longer required).



- Step 4. Add a `dls` local-peer command and one or more `dls` remote-peer commands (at the same remote site) that point to the central site DLSw+ routers.
- Step 5. At the central site RSRB device, remove the source-bridge remote-peer command that points to the device you just updated in Step 4.
- Step 6. Repeat Steps 3 through 5 with the remaining remote sites. When all remote sites are migrated, remove the local source-bridge remote-peer statement (and source-bridge `fst-peername` statement) from the central site router.

Caveats

Using a single router to perform both RSRB and DLSw+ during the migration phase is not always possible or advisable. In some situations, you might end up with loops, and in other situations you might be unable to establish SNA sessions

If both RSRB and DLSw+ remote peers need to access resources locally attached to the peering router via a bridge-group, you should not run DLSw+ and RSRB in the same router. Because of limitations of transparent bridging (that is, the lack of a RIF), you might incur problems if accessing the same bridge-group from both DLSw+ peers and RSRB peers (via the SR/TLB function).

If there are multiple source-bridge ring-groups defined on the RSRB peer device, great care must be taken when adding DLSw+. In particular, do not attempt to put DLSw+ on any RSRB peer that implements a de-encapsulation ring (a ring on which the RSRB peer has two or more interfaces, each bridging traffic to a separate ring-group as shown in Figure 8-1). Because of the way DLSw+ learns resource reachability, this type of scenario will cause problems.


If your RSRB router configuration contains multiple source-bridge ring-groups but does not include a de-encapsulation ring, then DLSw+ and RSRB can both run in the same router as long as you use ring lists to control which ring-groups DLSw+ uses. Ring-lists are required because unlike RSRB, all DLSw+ peers are tied to every ring-group defined in a router.

If the RSRB peer device in the data center is doing reverse SDLLC or reverse QLLC translation (providing upstream connectivity to an SDLC- or QLLC-attached device), the SDLC or QLLC link cannot be shared by RSRB and DLSw+. It is possible to share the router in these cases, but RSRB and DLSw+ must each have its own separate connection to the SNA device in question.

Migrating Any-to-Any Networks

To migrate an any-to-any RSRB network to DLSw+, it is best to use a de-encapsulation ring, which is a physical Token Ring that enables traffic flow between RSRB and DLSw+ networks. Starting with a fully meshed RSRB network, you move routers, one at a time, to either a fully meshed DLSw+ network, as shown in Figure 8-5, or to a peer group network. The two networks support any-to-any connectivity by using the de-encapsulation ring to move data between RSRB peers and DLSw+ peers.

Figure 8-5 Using a De-encapsulation Ring and Multiple Routers to Migrate Any-to-Any Networks



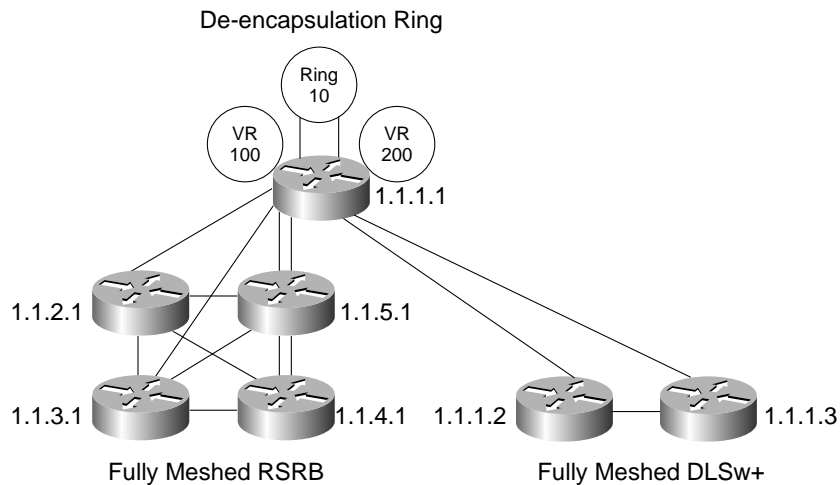
Because of the potential traffic volume on the de-encapsulation ring, this ring should be dedicated to this task. Do not use an existing ring that is handling other traffic. Additionally, during the migration there will be a heavy load on the routers connected to the de-encapsulation ring, especially where local acknowledgment is used. If there is not a lot of delay in the RSRB portion of the network, do not configure RSRB peers for local acknowledgment. This setup will minimize the overhead in the RSRB routers.

When using this type of migration, the DLSw+ peers on the de-encapsulation ring must use a different ring-group number than is used for the RSRB network. Otherwise, traffic will be unable to traverse the de-encapsulation ring from RSRB to DLSw+ or vice versa. As a result of this stipulation and the fact that DLSw+ terminates the RIF, it is important to maintain only one de-encapsulation ring, because explorers would otherwise jump from RSRB to DLSw+ and back again, increasing the overall explorer load on the network.

The de-encapsulation ring will be a single point of failure during the migration. To minimize the impact of this, you can use an intelligent hub that senses failures and wrap-around ports.

In Figure 8-5, a new router was added (attached to the de-encapsulation ring) for simplicity and to ensure that it could handle the load. Alternately, a single router attached to the de-encapsulation ring could be configured for both RSRB and DLSw+. A sample of this is shown in Figure 8-6.

Figure 8-6 Using a De-encapsulation Ring and a Single Router to Migrate Any-to-Any Networks



The following is a sample configuration for the router attached to the isolation ring in Figure 8-6:

```
source-bridge ring-group 100
source-bridge remote-peer 100 tcp 1.1.1.1
source-bridge remote-peer 100 tcp 1.1.2.1
source-bridge remote-peer 100 tcp 1.1.3.1
source-bridge remote-peer 100 tcp 1.1.4.1
source-bridge remote-peer 100 tcp 1.1.5.1

source-bridge ring-group 200
dlsw local-peer peer-id 1.1.1.1
dlsw port-list 1 TokenRing 1
dlsw remote-peer 1 tcp 1.1.1.2
dlsw remote-peer 1 tcp 1.1.1.3
interface TokenRing 0
 source-bridge 10 1 100
 source-bridge spanning
interface TokenRing 1
 source-bridge 10 1 200
 source-bridge spanning
```

Because RSRB remote peers are tied to a specific ring-group (in this case 100), only traffic from interface TokenRing 0 goes to RSRB peers, and traffic from RSRB peers only go out interface TokenRing 0. Because a port list is specified on all the DLSw+ peers, only traffic from interface TokenRing 1 goes to DLSw+ peers, and traffic from DLSw+ peers only go out interface TokenRing 1. If both of these interfaces are connected to the same physical ring, communication between the RSRB and DLSw+ domains is possible (as is desired during an any-to-any network migration).

Fully Meshed RSRB Network to Fully Meshed DLSw+ Network

If your current RSRB routers peer to every other RSRB router (that is, if you have a fully meshed RSRB network), this section describes the easiest way to migrate your network to a fully meshed DLSw+ network. However, you will have a better performing network if you use peer groups instead of a fully meshed DLSw+ network. You might decide to move to that type of design in the future, in which case this example can be viewed simply as the first step, where the second step adds DLSw+ border peers.

First, determine the location of the de-encapsulation ring. Your key goal is to minimize WAN impact. During the migration, all LLC2 sessions that need to traverse from an RSRB site to a DLSw+ site must pass through this de-encapsulation ring. Traffic might have to traverse the WAN into this site and then traverse it again to reach the destination site.

It is possible to use an existing router to attach to this new de-encapsulation ring if the device is being utilized lightly enough and can handle the extra load. Unless a network administrator is confident about the traffic patterns in the network, the ability of an existing router to handle the extra load can be very difficult to determine in advance. Observe the router during the course of the migration to determine when it is getting overloaded and needs some help. Important observations are CPU utilization, buffer utilization (tuning might be useful), and WAN link utilization to this site. Also observe the DLSw+ peers on this de-encapsulation ring, especially CPU utilization because the LLC2 session maintenance is a costly operation.


To migrate a fully meshed RSRB network to a fully meshed DLSw+ network using an existing RSRB router, do the following:

- Step 1. Ensure that your routers are at Cisco IOS Release 10.3(2) or later (see Appendix B to determine which version of Cisco IOS Software you should install to get the features you want).
- Step 2. Configure a `dlsw local-peer` command, specifying the `promiscuous` keyword, at the router attached to the de-encapsulation ring. If desired, at the end of the migration process the `static dlsw remote-peer` command can replace the `promiscuous` keyword. Also add a `dlsw port-list` command to keep RSRB and DLSw+ traffic separate.
- Step 3. Select an RSRB router to be migrated and delete the `source-bridge remote-peer` commands (and any related commands such as `SDLLC`, `SR/TLB`, `proxy explorer`, and `NetBIOS` name caching that are no longer required).
- Step 4. Add the appropriate `dlsw local-peer` command and a `dlsw remote-peer` command (at the same router) that is pointing to the DLSw peer on the de-encapsulation ring. On this `dlsw remote-peer` command, specify the same `port-list` number specified in Step 2.
- Step 5. Add a `dlsw remote-peer` command pointing to every router already converted to DLSw+.
- Step 6. Add a `dlsw remote-peer` command pointing to the router you just converted to DLSw+ to every router already converted to DLSw+.
- Step 7. Remove the `source-bridge remote-peer` command that referred to the router you just modified from all the RSRB routers, including the RSRB router on the de-encapsulation ring.
- Step 8. Repeat Steps 3 through 7 with the remaining RSRB routers.

You do not have to perform Step 7 immediately. RSRB is slightly degraded until it is done because RSRB attempts to connect to a peer that is no longer there. However, it will continue to work for all existing RSRB devices. Because Steps 6 and 7 require visiting the configuration of all devices in the network, you might migrate a batch of routers and do the corresponding updates for the entire batch at once.

To migrate a fully meshed RSRB network to a fully meshed DLSw+ network using a new router for DLSw+, do the following:

- Step 1. Ensure that your routers are at Cisco IOS Release 10.3(2) or later (see Appendix B to determine which version of Cisco IOS Software you should install to get the features you want).
- Step 2. Attach the new router and your co-located RSRB routers to the de-encapsulation ring. At the new router, configure a `dlsw local-peer` command with the `promiscuous` keyword specified. If desired, at the end of the migration process the `promiscuous` keyword can be replaced with `static dlsw remote-peer` command definitions.

- 
- Step 3. Select an RSRB router to be migrated and delete the source-bridge remote-peer commands (and any related commands such as SDLLC, SR/TLB, proxy explorer, and NetBIOS name caching that are no longer required).
 - Step 4. Add the appropriate dlsw local-peer command and a dlsw remote-peer command (at the same router) pointing to the dlsw peer on the de-encapsulation ring.
 - Step 5. Add a dlsw remote-peer command pointing to every router already converted to DLSw+.
 - Step 6. Add a dlsw remote-peer command pointing to the router you just converted to DLSw+ to every router already converted to DLSw+.
 - Step 7. Remove the source-bridge remote-peer command that referred to the router you just modified from all the RSRB routers, including the RSRB router on the de-encapsulation ring.
 - Step 8. Repeat Steps 3 through 7 for the remaining RSRB routers.

You do not have to perform Step 7 immediately. RSRB is slightly degraded until it is done because RSRB attempts to connect to a peer that is no longer there. However, it will continue to work for all existing RSRB devices. Because Steps 6 and 7 require visiting the configuration of all devices in the network, you might choose to migrate a batch of routers and do the corresponding updates for the entire batch at once.

At some point in the migration, you might find that either the one RSRB device or the one DLSw+ device on the de-encapsulation ring is insufficient to handle the load required of it. Because the RIF is not terminated in RSRB, it is safe to add additional RSRB peers to the de-encapsulation ring to help divide the load. Because all the RSRB peers within the same cloud share a ring-group number, it is clear that one peer will not read in frames originating from another RSRB peer; furthermore, no RSRB peer will put a frame onto the de-encapsulation ring that has already traversed it. Additional DLSw+ peers can be added to this ring to handle additional load as well; however, more caution must be used than in the RSRB case. DLSw+ permits different ring-group numbers to be used in the same DLSw+ cloud. It is important that all DLSw+ peers sharing this ring use the same ring-group number to avoid explorer looping. Also, because DLSw+ terminates the RIF, it is unable to determine whether an explorer it receives via its peers has traversed the ring already. Because of this situation, it is important that DLSw+ peers not peer directly to each other if they are both directly attached to the de-encapsulation ring.

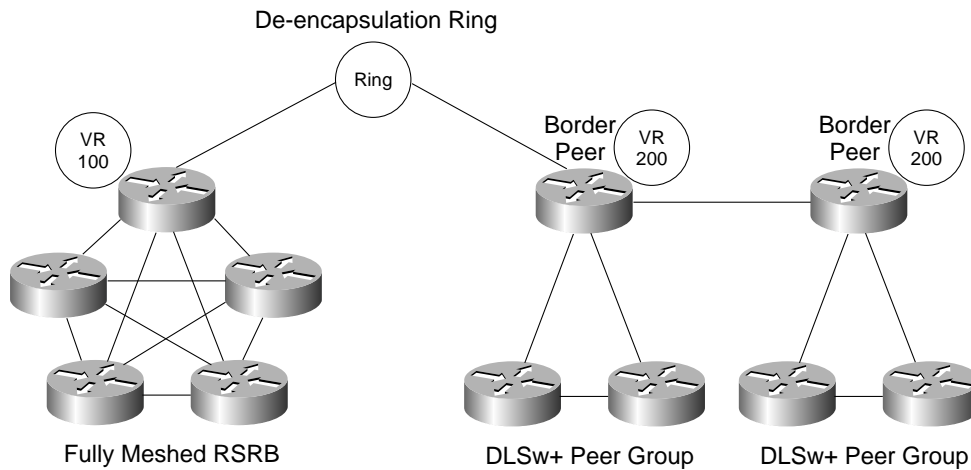
Before adding additional RSRB or DLSw+ peers on the de-encapsulation ring, it might be useful to examine the traffic patterns traversing the ring. It is possible that by identifying a particular RSRB site that is generating a lot of traffic to one of the DLSw+ sites that has already been migrated, the situation can be alleviated by migrating that RSRB site to DLSw+ to reduce this traffic flow. If possible, this setup would have the additional advantage of removing the requirement for multiple encapsulation steps for a large number of sessions, which should improve response time and overall network performance.

Fully Meshed RSRB Network to DLSw+ Peer Groups

Before beginning this migration, refer to the chapter “Designing Meshed Networks” to determine how you want to design the DLSw+ peer groups. After you have determined which routers will be border peers and which routers will belong to each peer group, you can begin the migration. Border peers should be migrated before any member peers in that group. You can migrate one group at a time or you can migrate all border peers first.

Many of the concepts discussed in the previous section still apply here. You still have a de-encapsulation ring as shown in Figure 8-7. Any DLSw+ peer on the de-encapsulation ring is configured as a border peer in its own group (there are no other peers in this peer group.) If additional DLSw+ peers are required to handle the load, then each one will be configured as a border peer in its own group (they will not share the same group number). Border peers on the de-encapsulation ring will not be peered to each other. (This rule is an exception to the general rule that all border peers in a DLSw+ network should be peered to each other.)

Figure 8-7 Migrating Fully Meshed RSRB to DLSw+ Peer Groups



This example uses new routers as DLSw+ border peers for the migration process and migrates one peer group at a time. To migrate a fully meshed RSRB network to a DLSw+ network with peer groups, do the following:

- Step 1. Ensure that your routers are at Cisco IOS Release 10.3(2) or later (see Appendix B to determine which version of Cisco IOS Software you should install to get the features you want).
- Step 2. Attach the new router and your co-located RSRB routers to the de-encapsulation ring. At the new router, configure a `dlsw local-peer` command with the `promiscuous`, `border`, and `group` keywords specified. If desired, at the end of the migration process the `promiscuous` keyword can be replaced with static `dlsw remote-peer` definitions.
- Step 3. Select an RSRB router to be migrated and delete the `source-bridge remote-peer` commands (and any related commands such as `SDLLC`, `SR/TLB`, `proxy explorer`, and `NetBIOS name caching` that are no longer required).
- Step 4. Add the appropriate `dlsw local-peer` command and a `dlsw remote-peer` command (at the same router) to point to the border peer configured in Step 2. The `group` keyword should be specified on the `dlsw remote-peer` command to indicate that this new peer belongs to the same peer group as the border peer to which it is being peered.
- Step 5. Remove the `source-bridge remote-peer` command that referred to the router you just modified from all the remaining RSRB routers, including the RSRB router on the de-encapsulation ring.
- Step 6. Repeat Steps 3 through 5 for the remaining RSRB routers that will belong to this peer group.
- Step 7. Add another border peer. This example assumes it is a new router. This router should not be attached to the de-encapsulation ring. At the new router, configure a `dlsw local-peer` command with the `promiscuous`, `border`, and `group` keywords specified. The `group` number configured for this new peer should differ from that used for other border peers already configured. If desired, at the end of the migration process the `promiscuous` keyword can be replaced with static `dlsw remote-peer` command definitions.
- Step 8. Configure a `dlsw remote-peer` command pointing to every other border peer already configured.
- Step 9. Add a `dlsw remote-peer` command pointing to this new border peer in every border peer already configured,
- Step 10. Repeat Steps 3 through 5 for all RSRB peers that will belong to the same peer group as this new border peer.
- Step 11. Repeat Steps 7 through 10 for all remaining peer groups and border peers.

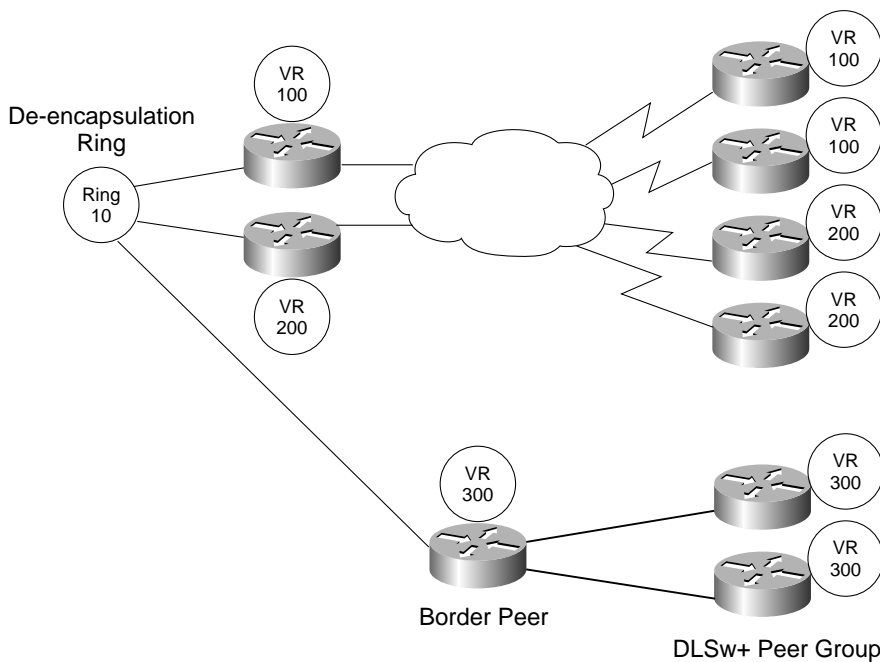
Note that for the majority of new DLSw+ peers (those that are not to be border peers), there are no configuration changes required for any other DLSw+ devices already converted (because the border peer in the group accepts the peer connection promiscuously). For the border peers, the number of previously converted DLSw+ peers that need to be touched is limited to a small subset (only the previously converted border peers).

Multihop RSRB to DLSw+ Peer Groups

For some RSRB environments that required any-to-any communication, full peer meshing was not possible because there were too many peers. For that reason, some large RSRB networks already use the concept of a de-encapsulation ring in order to connect two RSRB clouds using different ring-group numbers, as shown in Figure 8-8. When migrating to DLSw+ from this type of network, the key points to remember are as follows:

- The DLSw+ ring-group number should be different from any of the ring-group numbers being used for RSRB.
- Even if multiple de-encapsulation rings are being used, DLSw+ can connect into only one of them. RSRB supports multiple de-encapsulation rings because it does not terminate the RIF; DLSw+ cannot do this because of RIF termination (unless you are running Cisco IOS Release 12.0 and using RIF passthru).
- Migrate one RSRB ring-group before starting another. Hopefully the RSRB ring-groups were designed so that the majority of traffic flow remained within the same ring-group, with a minority having to traverse ring-groups. By migrating one ring-group before moving onto the next, less stress should be put on the routers attached to the de-encapsulation ring.

Figure 8-8 Migrating from Multihop RSRB to DLSw+ Peer Groups



This example uses new routers as DLSw+ border peers for the migration process and migrates one peer group at a time. To migrate a multihop RSRB network to a DLSw+ network with peer groups, do the following:

- Step 1. Ensure that your routers are at Cisco IOS Release 10.3(2) or later (see Appendix B to determine which version of Cisco IOS Software you should install to get the features you want).

- Step 2. Attach the new router and your co-located RSRB routers to the de-encapsulation ring. At the new router, configure a dlsw local-peer command with the promiscuous, border, and group keywords specified. If desired, at the end of the migration process the promiscuous keyword can be replaced with static dlsw remote-peer command definitions.
- Step 3. Select an RSRB router to be migrated and delete the source-bridge remote-peer commands (and any related commands such as SDLLC, SR/TLB, proxy explorer, and NetBIOS name caching that are no longer required).
- Step 4. Add the appropriate dlsw local-peer command and a dlsw remote-peer command (at the same router) that points to the border peer configured in Step 2. The group keyword should be specified on the dlsw remote-peer command to indicate that this new peer belongs to the same peer group as the border peer to which it is being peered.
- Step 5. Remove the source-bridge remote-peer command that referred to the router you just modified from all the remaining RSRB routers, including the RSRB router on the de-encapsulation ring.
- Step 6. Repeat Steps 3 through 5 for the remaining RSRB routers in the same ring-group.
- Step 7. Add another border peer. This example assumes it is a new router. This router should not be attached to the de-encapsulation ring. At the new router, configure a dlsw local-peer command with the promiscuous, border, and group keywords specified. The group number configured for this new peer should differ from that used for other border peers already configured. If desired, at the end of the migration process the promiscuous keyword can be replaced with static dlsw remote-peer command definitions.
- Step 8. Configure a dlsw remote-peer command pointing to every other border peer already configured.
- Step 9. Add a dlsw remote-peer command pointing to this new border peer in every border peer already configured.
- Step 10. Repeat Steps 3 through 6 for all RSRB peers that will belong to the same peer group as this new border peer.
- Step 11. Repeat Steps 7 through 10 for all remaining peer groups and border peers.

Multivendor Interoperability

You can configure a Cisco DLSw+ router to communicate with a non-Cisco router. However, not all the DLSw+ features will be available. Table 8-2 illustrates the features in DLSw+ that are beyond what the standard offers. Some of the DLSw+ features can be used (with some restrictions) even if the peer at the other end is not a Cisco router.

This section details the options you cannot configure and the options that are configurable but somewhat unpredictable. All interoperability testing was done at base-RFC 1795 and base-RFC 2166 levels only. Cisco has tested interoperability with several vendors. Contact Cisco to find out the latest status of this interoperability testing.

Table 8-2 Comparison of DLSw+ and Standard DLSw Features

	DLSw Standard Feature	Additional DLSw+ Features
Performance	IP load sharing Circuit-level flow control	Peer ¹ , port and RIF load sharing Custom and priority queuing Weighted fair queuing and weighted random early detection ToS/COS mapping RSVP support FST encapsulation



Table 8-2 Comparison of DLSw+ and Standard DLSw Features (Continued)

Availability	Nondisruptive rerouting No data-link control timeouts	Backup peers ¹ Fault tolerant peers ¹ Load sharing across peers
Scalability	Broadcast reduction Hop count reduction UI/UDP support Multicast	Broadcast optimization with peer groups Border peer caching Ring lists ¹
Flexibility	Media conversion SRB dynamics Capabilities exchange for cache preloading	Peer biasing with cost ¹ SNA DDR Diverse data-link control media (QLLC, Reverse SDLLC, Token Ring LANE, Token Ring ISL, SRB over FDDI) Media conversion between SDLC and LLC2 for PU 4-to-PU 4 Dynamic peers DLSw Lite

1. Can be used with a non-Cisco router

The key limitations when building networks with a mix of DLSw standard and DLSw+ routers are as follows:

- You cannot use any encapsulation other than TCP (to non-Cisco routers).
- Non-Cisco routers cannot be border peers or participate in peer groups; in a mixed-vendor environment, you also might not be able to take advantage of load balancing, backup peers, and cost (for these features, it depends on which router is the Cisco router and which one is the non-Cisco router).
- A Cisco router load balances between two central site non-Cisco routers as long as the Cisco router initiates the CANUREACH exchange; a Cisco router also locally load balances across all interfaces.
- Cisco routers automatically connect to a backup peer upon loss of a primary peer even if the backup peer is a non-Cisco router. The non-Cisco router must either be able to accept a connection from an unknown peer or support something equivalent to the passive keyword. The Cisco router automatically terminates the backup peer connection according to the configuration options.
- Cisco routers establish a dynamic peer connection with a remote non-Cisco peer as long as that remote peer accepts either a connection from an unknown peer or support something equivalent to the passive keyword.
- Cisco routers bias remote peer selection based on cost and can support diverse local media.

Again, interoperability testing has been limited to RFC 1795 features only, but all of the features should work. Most of these features require that the Cisco router be at the initiating end of the connection.

Local Peer Statements

The following `dlsw local-peer` command keywords are valid because they do not contain information that is sent in a capabilities exchange to a remote router or are specified in the RFC:

```
dlsw local-peer [peer-id ip-address] [lf size] [keepalive seconds] [passive] [promiscuous] [biu-segment]
[init-pacing-window size] [max-pacing-window size]
```

The following `dlsw local-peer` command keywords can be configured in a Cisco router, but they should be ignored by non-Cisco, standard-compliant routers (they are passed in the capabilities exchange as vendor-specific vectors):

```
dlsw local-peer [group group] [border] [cost cost]
```

Remote Peer Statements

For `dlsw remote-peer` commands, you must specify TCP/IP encapsulation. The following keywords are available on this command:

```
dlsw remote-peer list-number tcp ip-address [backup-peer ip-address] [bytes-netbios-out bytes-list-name] [cost cost] [dest-mac mac-address] [dmac-output-list access-list-number] [host-netbios-out host-list-name] [lf size] [linger minutes] [lsap-output-list list] [tcp-queue-max size]
```

These keywords control local filtering, biasing, queue depths, and control when this peer will initiate disconnects with a remote peer.

You should not use the following keywords when you configure the dlsw remote-peer commands:

```
dlsw remote-peer list-number tcp ip-address [dynamic] [inactivity minutes] [keepalive seconds] [no-llc minutes] [priority] [timeout seconds]
```

The priority keyword should not be configured, because it causes a Cisco router to open four TCP queues, which another vendor's router might not understand or accept. Whether the dynamic keyword works as desired is vendor-dependent. Associated with the dynamic keyword are inactivity *minutes* and no-llc *minutes*. SNA DDR relies on timeout seconds to control when TCP recognizes that it has lost a connection and keepalive seconds to be set to zero to prevent keepalives from keeping up dial lines. Both keywords have unpredictable results when used with another vendor's router.

Other DLSw+ Commands

Other DLSw+ configuration commands that can be used include:

- dlsw ring-list
- dlsw mac-address
- dlsw netbios-name
- dlsw icanreach
- dlsw load-balance