

Designing Hierarchical Networks

Hierarchical DLSw+ networks are the easiest networks to design and build. They involve minimal routing and are inherently scalable. If you are going to design a hierarchical DLSw+ network, you must answer the following questions:

- How many central site routers are required to handle the traffic load?
- Where is the best place for the central site peer routers?
- How will backup be performed?
- What can be done to minimize explorer traffic and broadcast replication?

This chapter discusses each of these questions and provides information to assist you in making the best decisions for your network. Read this chapter if you are connecting several remote branches to a single primary data center. You may also need to read the “Designing Meshed Networks” chapter if you have frequent branch-to-branch communication among SNA or NetBIOS applications.

Determining the Required Number of Peering Routers

There are many factors involved in determining the number of central site routers required to support a hierarchical network. These factors include the following:

- Number of SNA PUs or concurrent LLC2s to be supported
- Transaction rate at central site and transaction size
- Encapsulation method selected
- Central site routers used for peering
- Number of remote peers connected
- Explorer replication
- Other router processes, such as multiprotocol routing and route table maintenance, compression, and encryption

Number of Devices

The number of SNA PUs is relevant when local acknowledgment is used because each SNA PU has an SDLC or LLC2 connection that must be kept alive by sending messages at regular intervals. These keepalive messages and the timer processing required to determine when to send them is processor intensive. Adjusting LLC2 timers on the routers can help, but in general, on a Cisco 4700 Series router assume a maximum of approximately 4000 PUs. Figure 6-1 and Figure 6-2 illustrate the CPU utilization of various routers for varying numbers of PUs and LUs and can be used to approximate the size of the router required. For a more exact calculation, provide the appropriate information to your systems engineer.

Figure 6-1 CPU Usage of Various Routers Assuming TCP Encapsulation and Assuming Each PU Has 10 LUs, Each with One Transaction per Minute

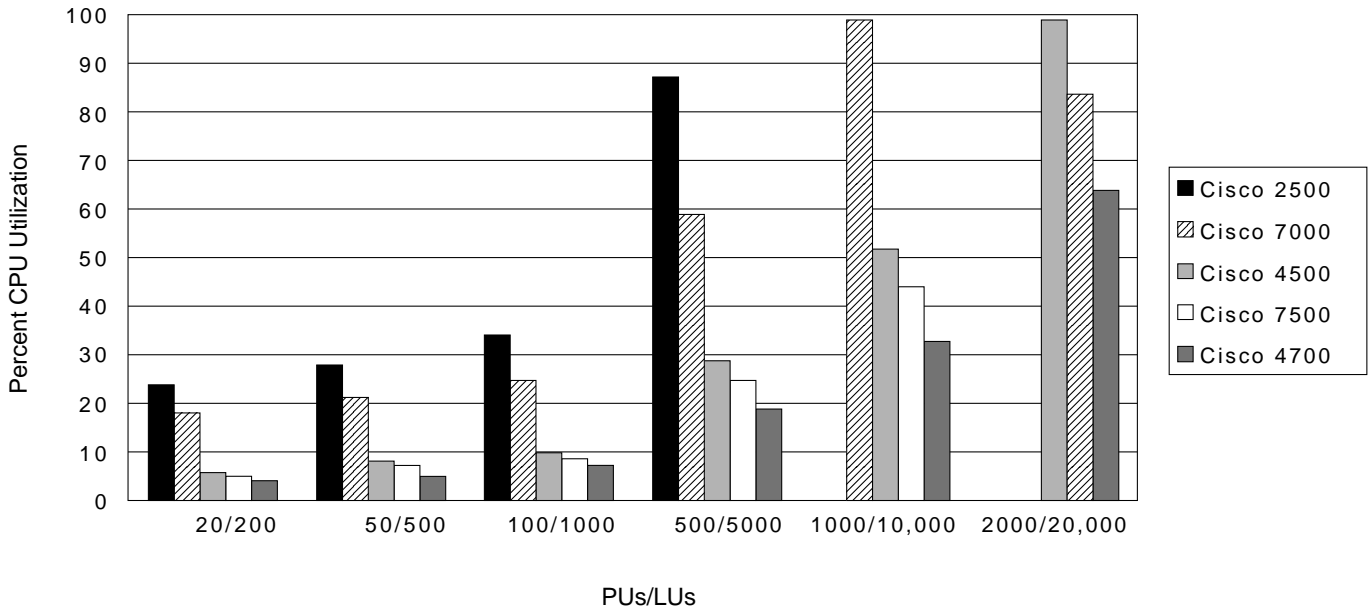




Figure 6-2 CPU Usage of Various Routers Assuming TCP Encapsulation, Transactions of 40 Bytes in and 1000 Bytes Out, and Assuming Each PU has 4 LUs, Each with One Transaction per Minute



Transaction Rate

The transaction rate also plays a role in determining how many central site routers can be supported. A typical transaction rate is one transaction per LU per minute. By determining the number of LUs per PU on average and the total number of PUs and assuming this transaction rate, you can fairly accurately anticipate the transaction rate of most environments. The transaction size has two components: message size in and message size out (40 bytes in and 1000 bytes out is common). Figure 6-1 and Figure 6-2 illustrate the router utilization with a specific



transaction rate and size. Note that the number of PUs has more of an impact than the transaction rate. Varying the LUs is relevant because it changes the transaction rate. If the transaction rate was kept constant as new LUs were added (in other words, fewer transactions per LU as LUs were added), the number of LUs would have no bearing.

Encapsulation Method

The encapsulation method is relevant because different encapsulation methods have different impacts on route processor utilization. Both TCP and LLC2 encapsulation involve local termination of the data-link controls (local acknowledgment) and are process switched. FST and direct encapsulation run in passthru mode, which means acknowledgments flow end to end. Assuming adequate bandwidth and line quality, these encapsulation types will allow a central site router to support more remote branch routers, because these encapsulations do not support local acknowledgment and require fewer processor cycles. Figure 6-1 and Figure 6-2 assume TCP encapsulation.

Processor Speed

DLSw+ is processor intensive and runs best in a router with a faster route/switch processor (for example, a Cisco 4700, 7200, or 7500 Series router) rather than a slower processor (for example, a Cisco 4000 or 7000 Series router). Figure 6-1 and Figure 6-2 show the CPU utilization required to support various numbers of PUs and traffic volumes. In Figure 6-1, the transaction size was 40 bytes in and 1000 bytes out. Each PU had 10 LUs, and each LU transmitted at a rate of one transaction per minute. Using these numbers, 500 PUs and 5000 LUs result in 5000/60, or 83 transactions per second at the central site router. The LLC2 idle timer on the Token Ring interface was set to 30 seconds for these tests.

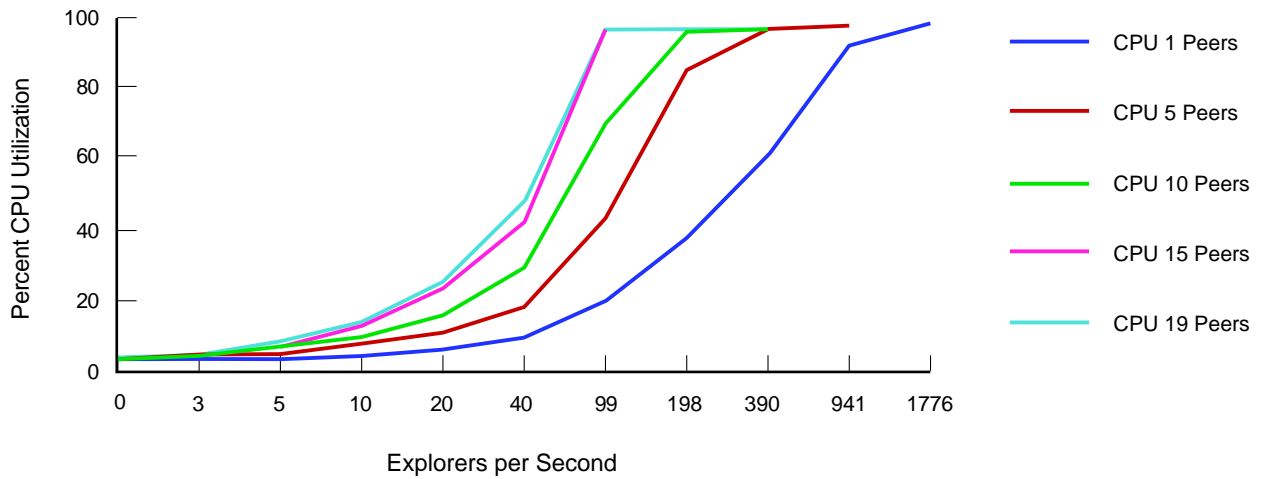
Number of Remote Sites

The number of remote sites that must be connected may have an impact on the number of central site routers required. This number is important when broadcasts must be replicated (see the section “Explorer Replication”). Performance testing shows that the number of peers has no significant impact on router CPU usage if there is no broadcast traffic.

Explorer Replication

Another factor that plays a role in determining the number of central site routers is the amount of explorer replication required. If all the connection requests are remotely initiated and the network is hierarchical, the amount of explorer replication required should be minimal. This assumes appropriate filters are set at the central site to prevent unnecessary explorer propagation (see the “Customization” chapter). When connection requests are initiated at a central site, these requests must be propagated to each remote peer. The number of explorers that can be replicated per second depends on the speed of the route processor. Figure 6-3 illustrates the explorer processing rate of a Cisco 4700 Series router. As the number of peers increases, the number of explorers that can be received and replicated per second decreases. For example, if a Cisco 4700 Series router peers to 20 remote peers, it can replicate almost 100 explorers per second to each of the 20 peers. If the router peers to one router, it can replicate more than 1700 explorers.

Figure 6-3 CPU Utilization of a Central Site Router as the Number of Explorers per Second Varies and as the Number of Peering Routers Increases (No Caching Is Assumed, and Each Explorer Received Must Be Replicated to Each Remote Peer)



Other Router Processes

Your router may be configured to do more than DLSw+. The NetSys Performance Solver tool will help you size routers that are performing multiple functions, or you can approximate the number of routers required based on the amount of CPU you are using for routing functions and what additional load your SNA traffic will place on those routers.

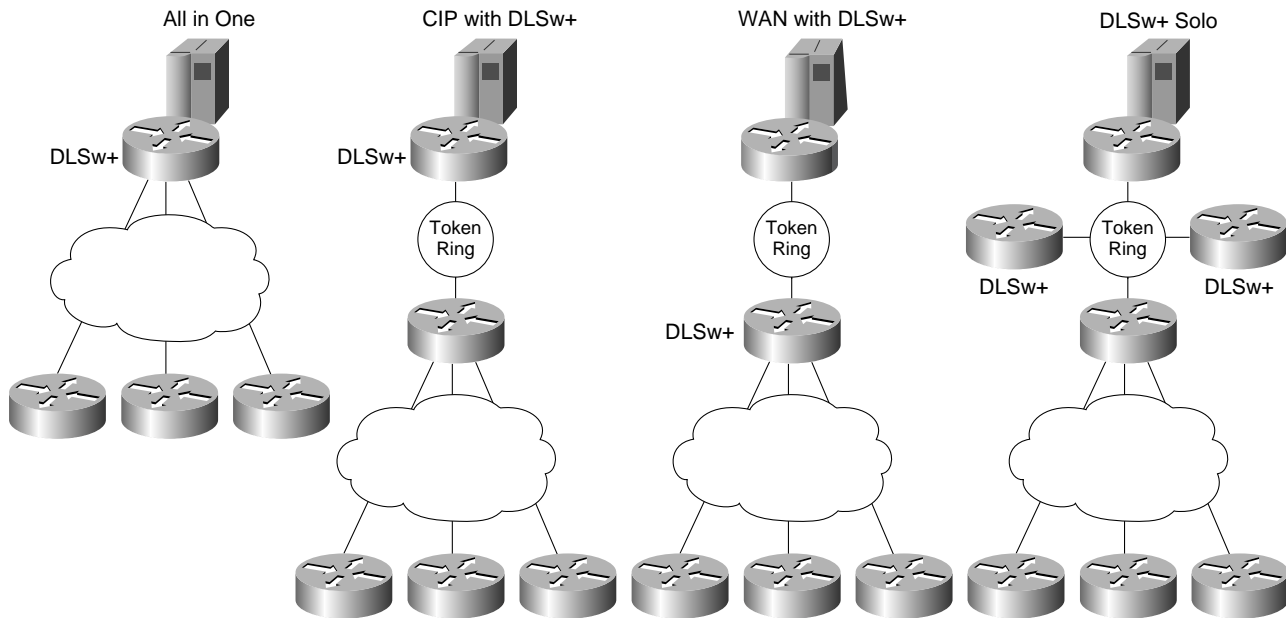
Placement of Peering Routers

After you have determined how many routers you need to support your traffic, you can consider the best place to put the peering routers. There are four alternatives:

1. Peer all remote sites to one or more central site routers that are directly connected to a mainframe over a CIP and directly connected to the WAN over serial ports
2. Peer all remote sites to one or more central site routers that are directly connected to a mainframe over a CIP, but keep WAN function in separate routers
3. Peer all remote sites to direct WAN-attached routers that access the mainframe via a channel gateway such as an IBM 3745 or another Cisco router with a CIP
4. Peer all remote sites to dedicated DLSw+ routers that are neither WAN connected nor CIP connected

Each of these alternatives is valid and is the best alternative in specific environments. Figure 6-4 illustrates these alternatives.

Figure 6-4 Four Central Site Peering Router Replacement Alternatives



All in One (DLSw+, CIP, and WAN)

Peering to a CIP router that is also a WAN router has the advantage that it requires the smallest number of central site routers. In small networks (30 to 60 branches) that are primarily SNA, this is a reasonable choice.

CIP and DLSw+ Combined

Peering to a CIP router but having a separate WAN router is a good solution for small to medium-sized networks (up to 200 remote branches) with a moderate amount of multiprotocol traffic. This design allows you to segregate multiprotocol broadcast replication from DLSw+ processing. For backup and availability, this solution will typically involve two central site peering routers; one router should be able to handle the load in the event of a failure of the other router.

WAN and DLSw+ Combined

The third solution, peering to the WAN router, is a good solution for medium-sized to large networks that require more than one or two central site routers for DLSw+ processing or that use a channel gateway other than a CIP. To access the channel gateway, you can use SRB over Token Ring (this is adequate for most host traffic) or SRB over FDDI (DLSw+ will support SRB over FDDI in Cisco IOS Release 11.2).

Using WAN and DLSw+ combined, you can segregate the DLSw+ processing from the CIP-attached router and scale the network without buying additional CIP routers. As the network grows beyond the capacity of a single router, you can add Cisco 4700 or 7200 Series routers to handle the capacity. This is more cost effective than adding large Cisco 7500 Series routers with CIPs. Because SRB is fast switched, a single Cisco 7500 or 7000 Series router with a CIP can handle the traffic from four or five Cisco 4500s or three Cisco 4700s using DLSw+. Figure 6-4 illustrates the transaction processing power of a Cisco 7500 and a CIP when using SRB to send traffic from a LAN to the CIP. The message size is noted in the first column, and the number of SNA PUs is indicated in the LLC2 column. The packets per second in and out and thousand bits per second in and out is shown in the

next two columns. The Cisco 7500 Route Switch Processor (RSP) utilization is always relatively low, because the traffic is fast switched off the LAN and to the CIP. The CIP is designed to run at 100 percent for some traffic volume.

To put these traffic volumes in perspective, a transaction rate of approximately 1300 per second would represent the load of 78,000 LUs sending data at a rate of one transaction per LU per minute. The RSP in this example was 26 percent utilized and the CIP was operating at 100 percent. These tests were run using a CIP1. All new CIPs are CIP2 and they support a much higher transaction rate.

This configuration also offers advantages in terms of change management and network availability. By limiting the channel-attached routers to SRB and IP routing, you minimize the requirement for configuration changes or Cisco IOS Software upgrades in your channel-attached router. This configuration decreases planned downtime and increases network reliability.

Dedicated DLSw+

The final alternative separates DLSw+ processing from CIP processing and WAN processing. This is a good solution for large networks with a significant amount of multiprotocol traffic. Although this appears to have the most routers, it may in fact have the same number of routers with the function split across different boxes. The key advantage to this solution is load balancing and backup. If the WAN is a Frame Relay network, a single permanent virtual circuit (PVC) to a central site WAN router provides connectivity to multiple central site peering routers. This configuration has the same change management and availability advantages as the previous one.

Note: FST or TCP encapsulation is required whenever the peering routers are not adjacent, as shown in the CIP with DLSw+ or DLSw+ Solo solutions in Figure 6-4. DLSw Lite and direct encapsulation options assume that the peering routers are adjacent (that is, that DLSw+ is running in the WAN router).

Availability Options

There are several alternatives for building a fault-tolerant network. With DLSw+, recovery from some failures is nondisruptive to the end systems. Recovery from any failure can be dynamic. The following describes recovery scenarios with various features.

Link Recovery

Link failures on the WAN can be recovered by using TCP encapsulation and providing alternate paths (either leased or switched). Local acknowledgment ensures that the router has time to reroute around the link failure without disrupting SNA sessions. Some NetBIOS applications have session-level timers in addition to link-level timers. DLSw+ does not spoof session-level timers, so NetBIOS sessions may drop if there is an outage in the network.

When using FST encapsulation, link failures may or may not be disruptive. Because FST does not offer local acknowledgment, timers may expire before DLSw+ has time to reroute. However, rerouting is dynamic.

When using direct encapsulation, link failures are disruptive but recovery can be automatic. Backup from link failures can be addressed either by having multiple remote peers or by configuring multiple remote peer statements for the same remote peer but specifying a unique path to each one. You can either load balance between them or use cost to cause one path or peer to be preferred over the other, as shown in the following statements:


```
dls w remote-peer 0 frame-relay interface serial 0 22 cost 2
dls w remote-peer 0 33.33.33.33 cost 4
```



In this example, the first statement describes how to get to a remote peer directly over a Frame Relay link, and the second statement describes how to get to the same remote peer via a TCP path.

Recovery using two peers is illustrated in Figure 6-5.

Figure 6-5 DLSw Lite Configuration Providing Dynamic Recovery from the Loss of a Link or Central Site Router



```
Configuration for Router A
dlsw local-peer peer-id 10.2.17.1
dlsw remote-peer 0 interface serial 1 33 cost 1
dlsw remote-peer 0 tcp 10.2.18.2 cost 4
interface serial 1
encapsulation frame-relay
frame-relay map llc2 33
```

```
Configuration for Router B
dlsw local-peer peer-id 10.2.24.3
promiscuous
duplicate-path-bias load-balance
interface serial 1
encapsulation frame-relay
frame-relay map llc2 17
```

```
Configuration for Router C
dlsw local-peer peer-id 10.2.18.2
promiscuous
duplicate-path-bias load-balance
```

Central Site Router Recovery

Loss of a central site router is always disruptive, but recovery can be dynamic and immediate. There are two alternatives for recovery: multiple concurrently active central site routers or backup peers.

If remote peers concurrently connect to multiple central site peers, loss of a single central site peer causes all new sessions to be established over the remaining active central site peers. If remote peers only connect to a single central site peer, you can still specify a backup peer that will be used in the event that the primary peer goes away.

See the “Advanced Features” chapter for a description of these features and a comparison of the alternatives.

Central Site Mainframe Channel Gateway Recovery

Loss of a central site mainframe channel gateway is always disruptive, but recovery can be dynamic and immediate. The simplest way to recover from this is to have multiple mainframe channel gateways with the same MAC address, each accessible via a different port on a central site router. DLSw+ supports load balancing for up to four ports. This not only addresses availability, it can also spread the traffic across multiple TICs on a FEP to avoid congestion problems. This configuration is commonly known as the duplicate TIC configuration. The same concept can be used in conjunction with a Cisco CIP. DLSw+ allows remote SDLC or Ethernet-attached devices to take advantage of this feature by providing media conversion.

Broadcast Reduction

Because broadcasts impact the processing power of a DLSw+ router, it is important to understand how to eliminate any unnecessary broadcast replication. Some techniques to eliminate replication are filtering, virtual ring numbering, and static device configuration.

Filtering

Filtering unnecessary broadcasts is the best way to minimize explorer replication. DLSw+ attempts to switch nonrouted multiprotocol traffic if not filtered. The “Customization” chapter describes how to configure filtering to allow only SNA or NetBIOS into DLSw+. When an access list has been created, it can be applied to the input interface (which would prevent even local forwarding) or to the `dlsw remote-peer` command.

Virtual Ring Numbering

When there are multiple central site DLSw+ routers attached to the same Token Ring segment or SRB LAN, it is possible for broadcast frames to come in from the WAN, be sent out over the physical Token Ring LAN, and be picked up by another DLSw+ router. To prevent that router from retransmitting the frame on the WAN, code the same virtual ring number in all DLSw+ routers attached to the same physical ring or bridged LAN. Normal SRB procedures prevent broadcasts from being copied on a ring that is already present in the RIF.

Static Device Configuration

Devices can be statically configured in DLSw+. By configuring frequently accessed resources, you can eliminate the need for broadcasts to find those resources.

DLSw+ allows you to statically configure resources (MAC addresses or NetBIOS names) that are local to a DLSw+ peer using a `dlsw icanreach` command. This information is dynamically distributed to all remote peers as part of the capabilities exchange. This feature is extremely useful as a means to advertise reachability of the mainframe channel gateway (FEP or CIP) or key NetBIOS servers. With a few configuration statements at central site routers, you can preload the cache of all the remote peers. When a peer learns of the reachability of an end system via a capabilities exchange, it keeps that information in its cache as long as the peer connection is active. The peer never broadcasts explorers for these resources. If a branch router peers to multiple central site routers, it can learn of multiple ways to access a resource and will cache up to four of them.

Central site routers can also specify the exclusive keyword. For example, the commands `dlsw icanreach mac-address 4000.3745.0001` and `dlsw icanreach mac-exclusive` tell remote routers that the *only* destination this router can reach is the MAC address of the IBM 3745. This feature can also be used to indicate that certain NetBIOS servers are located at the central site, but not elsewhere. The exclusive keyword prevents remote sites from forwarding unnecessary broadcasts.

DLSw+ allows a peer router to advertise when it cannot reach a resource or SAP (this is specified in the `dlsw icanotreach saps` command). One use of this feature is to prevent branch offices from searching the data center for NetBIOS servers. If a DLSw+ peer learns via a capabilities exchange that it cannot reach a resource via a particular peer, it will not send explorers to that peer for that resource.

Note: If you are using border peers there are some limitations. Because border peers offer no advantages for hierarchical networks, this chapter assumes they are not being used. See the “Designing Meshed Networks” chapter for a discussion on the implications of border peers and using `dlsw icanreach` configuration commands.

DLSw+ also allows you to statically configure a path to reach a local or remote resource. This is done using a `dlsw mac-addr` or `dlsw netbios-name` command, and it works well if there is only one way to reach a resource and its location will never change. This entry is never deleted from the cache. The “Customization” chapter describes the difference between using static paths and `dlsw icanreach` commands.