

# Designing Meshed Networks

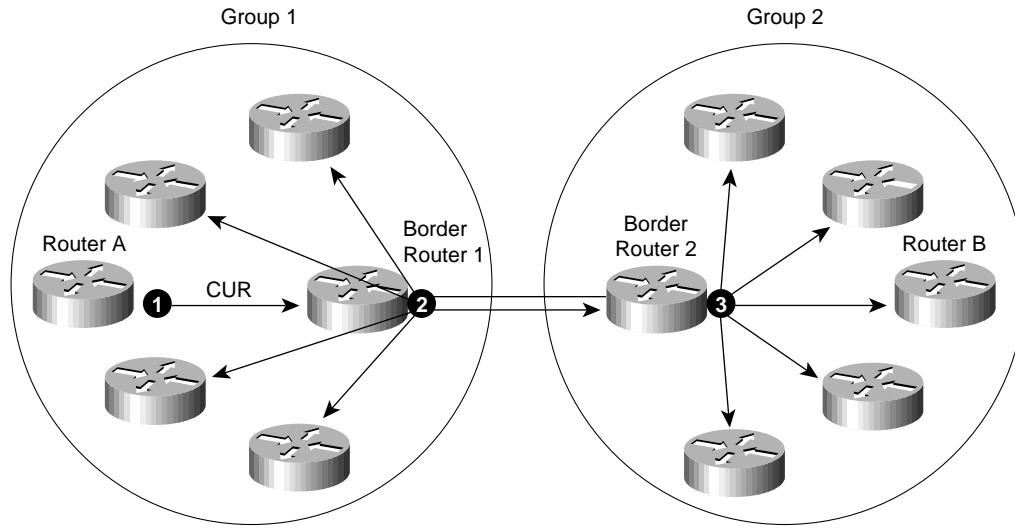
This chapter describes design considerations when using DLSw+ to build a meshed, any-to-any network. It describes the role of border peers and provides guidance on where to place border peers, how to backup border peers, and how to size border peers. It also describes how to size peer groups, how to configure on-demand peers, and how to minimize unnecessary broadcast replication.

## The Peer Group Concept

DLSw+ offers features designed specifically to address the issues in building fully meshed networks. These features are border peers, peer groups, and on-demand peers, and they are described in the “Introduction” and “Advanced Features” chapters. Using peer groups you can build enterprise networks that support branch-to-branch communication between either NetBIOS or Advanced Program-to-Program Communication (APPC) applications. Most enterprises do not require fully meshed connectivity, but when it is required, it can be challenging to support. The key reason is that unless resources are statically configured everywhere, extensive broadcast traffic results. This broadcast traffic can clog access links, and broadcast replication can overburden branch routers.

DLSw+ solves this problem by providing a hierarchical means to dynamically search for branch resources. Instead of a single branch router having to query every other branch router, the branch router sends a single broadcast to its border peer. The border peer checks its local, remote and group cache before forwarding the explorer. If the resource is not found, the border peer propagates the broadcast within its group and to other border peers. Other border peers propagate the broadcast within their group. This method not only minimizes the broadcast replication on each line, it also minimizes the replication work done by any single router. End-to-end TCP connections (called peer-on-demand connections) are set up only when resources are found. Figure 7-1 illustrates explorer processing when border peers are used.

Figure 7-1 Explorer Frames Are Processed by Border Peers




In Figure 7-1, Router A sends a single CANUREACH frame to Border Router 1. Border Router 1 checks its local, remote and group cache. If the resource is found in its cache, the border peer forwards the explorer to the known destination. If the resource is not found, then it forwards this broadcast to the remaining four routers in Group 1 and to Border Router 2. Border Router 2 forwards the broadcast to all the routers in Group 2. In this way all 10 branch routers receive the CANUREACH frame, but the originating branch router sent only a single copy, and each border router replicated it only five times.

DLSw+ further reduces explorer replication with the Peer Group Cluster feature, introduced in Cisco IOS Release 12.0(3)T. If multiple routers are serving the same LAN within the same peer group, they can be “clustered.” This classification further reduces explorer replication because the border peer does not send an explorer to more than one router serving the same LAN within a peer group.

## Explorer Processing When Using Peer Groups

When a `dlsw local-peer` command specifies a group number, that local peer forwards explorer packets only to border peers in its group, configured remote peers that are not part of any group (for example, non-Cisco routers), or peers in its local cache that match the explorer conditions. You can specifically configure remote peers within the same group, but the local peer still sends explorers only to its border peer. The only exception is that if there are no active border peers within its group, a local peer forwards broadcasts to all configured peers. For example, in Figure 7-1 if Router A is peered to all the routers in Group 1 and to Router B, when Router A gets an explorer frame, it forwards it only to Border Router 1. If Border Router 1 is not active, it forwards the explorer to all the peers in Group 1 in addition to Router B.

There are two reasons you may wish to configure all remote peers within a group. Again, using Figure 7-1 as an example, if all the routers in Group 1 frequently communicate with each other, by configuring `dlsw remote-peer` commands between every pair of routers within the group, the peer connections are always active. This shortens the connection time for end-user sessions without increasing the broadcast traffic. Connection to less frequently accessed peers in other groups can still be made using on-demand peers. In addition, if there is only a single border peer, it eliminates the single point of failure condition for connectivity within the group.



DLSw+ does not support cascaded groups. That is, if a border peer from one group forwards an explorer to a border peer in another group, that border peer does not forward the explorer to a third group.

Border peers will forward explorers to local interfaces in addition to other member peers.

## Border Peers

Currently, the sole function of border peers is broadcast replication on behalf of branch routers (or member peers). By concentrating this function in a more powerful distribution or in central site routers, and by distributing the replication function across a number of border peers, you can build networks with fully meshed connectivity without having to put large, powerful routers at every branch. The border peer should be a Cisco 4700, 7200, or 7500 Series router. The placement of the router is determined by your physical network design. It can reside either at a distribution site or at a central site, depending on where you send your branch traffic enroute to other branch sites.

DLSw+ supports multiple active border peers. Every peer in a group will forward explorers to one of the border peers in its group. If the border peers are configured with different costs, then the member peer will select the border peer with the lowest cost. If the border peers are configured with same cost, then the member peer will select the border peer with which it had the most recent active peer connection. Border peers do not perform load balancing.

If you are using multiple active border peers, the following rules apply:

- Within a single group, every member peer must peer to every border peer in its group
- All border peers within a group must peer to each other
- All border peers within a group must peer to every border peer in other groups
- Border peers forward explorers to all member peers in their group, all border peers in their group, and to one border peer in every other group


## On-Demand Peers

With border peers in place, it is possible for two peers to communicate with each other even though neither has a configuration for the other. This is because they learn about each other via their respective border peers. For that reason, there is a statement that defines how to connect with peers when no dlsw remote-peer commands are used. That statement is the dlsw peer-on-demand- defaults tcp command, and it can be used to specify the encapsulation, filters, and timers. On-demand peer connections can be made between two peers in the same group or two peers in different groups.

## Size of Peer Groups

How you divide your network into groups determines how much broadcast replication any single router must do. For example, with a 50-branch network, it is possible to use a single peer group as long as the broadcast traffic is not excessive. With a 1000-branch network, a single peer group is not practical (a single broadcast would have to be replicated 999 times). Suppose you designed a network with four peer groups, each peer group consisting of 250 routers. With this design, the border peer must replicate every broadcast once for every peer in its group, and once for every other border peer. As shown in Figure 7-2, with four groups of 250 routers, there are  $3 + 249$ , or 252, replications per broadcast. With 20 groups of 50 routers, there are  $19 + 49$ , or 68, replications per broadcast. You should design your groups in a manner that ensures your border peers can handle the amount of broadcast replication any single router must perform.

Figure 7-2 Number of Replications per Broadcast that a Border Peer Must Perform Based on Different Means of Splitting 1000 Branch Routers into Peer Groups




## Broadcast Replication/Reduction

Figure 7-3 shows the router utilization based on the number of explorers per second. The practical limit for a Cisco 4700 Series router is 800 to 1000 explorers per second, although it can replicate around 1800 explorers per second if it does nothing else. Assuming a forwarding limit of 800 to 1000 explorers per second, if the router has 20 peers, it can handle an incoming rate of 40 to 50 explorers per second.

Assume a border peer has 50 peer connections. The worst case scenario for broadcast replication is when a key resource is lost. Every end system tries to find that resource, resulting in at least 50 simultaneous broadcasts (each DLSw+ peer may get multiple explorers but forwards only the first request and queue any duplicates). This would require  $50 \times 49$ , or about 2500, explorers. The way to avoid this situation is to preconfigure any resources that an entire enterprise needs to access frequently. Border peers should only be used to find resources that are accessed on an as-needed basis and are not accessed by all branches, all the time.

In Cisco IOS Release 11.3, border peers were enhanced to support group caching, which greatly reduces explorers. If you have an earlier version of software, and you have a hierarchical SNA network and a meshed NetBIOS network, you need to consider the impact of border peers on explorer traffic. You may have occasional branch-to-branch traffic, but the resources that every branch accesses every day are at the central site (the FEP or enterprise servers). Every time a branch router needs to search for the FEP (because the MAC address of the FEP is not in its cache), the branch router sends an explorer to its border peer. The border peer forwards the explorer to every router in its group and every other border peer. Even if a border peer has found the FEP on behalf of another resource, it forwards the explorer everywhere. That is because, prior to Cisco IOS Release 11.3, border peers do not check their cache before forwarding explorers.

Figure 7-3 CPU Utilization Comparison 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)



As illustrated in Figure 7-4, to avoid unnecessary broadcast forwarding for central site resources, you can configure all the remote branch routers to peer to both their border peer and a data center router (Router C). At Router C, you can configure the reachability of the FEP MAC address in a `dls w icanreach` command. As soon as the branch router establishes a peer connection with Router C, it learns that Router C can access the MAC address of the FEP. When Router A receives an explorer for the FEP MAC address, it first checks its cache, where it finds a match (remember, cache entries learned as part of a capabilities exchange are not deleted unless the associated peer connection goes away). Instead of forwarding the explorer to its border peer, it forwards the circuit setup request directly to Router C, avoiding any broadcasts to other branch routers.

**Note:** In this replication scenario, if the border peer and the data center router are the same, specify `dls w icanreach` in the border peer. In this way, the branch router preloads its cache with the MAC address of the FEP and always sends a directed explorer to the border peer instead of requesting that the border peer do a search on its behalf.

Figure 7-4 Explorers Are Minimized by Peering Branch Routers to Both Their Border Peer and a Central Site Router

