# Optical Networking Systems IP Management Solutions

Randy Zhang

ciscopress.com



shortcut

Your Short Cut to Knowledge

# About the Author

***Randy Zhang***, Ph.D. and CCIE No. 5659, is a network consulting engineer at Cisco. He has authored many white papers on subjects relating to IP and optical technologies and spoke on optical timing and ONS IP management at Cisco Networkers conferences. His white paper on timing and synchronization has widespread readership. Randy is a co-author of the Cisco Press book *BGP Design and Implementation*, published in 2003.

# About the Technical Editor

***Lee Shombert***, Ph.D., is a software engineer currently with the Cisco Gigabit Switching Business Unit. Prior to his position at GSBU, he was in the Cisco Optical Networking Group, where he designed and developed many of the IP networking features on the ONS 15454 platform, including the firewall, SOCKS proxy, and craft access features.

## Dedication

To Susan, Amy, and Ally

## Acknowledgements

I would like to express my deep gratitude to Lee Shombert, my technical reviewer. Lee's technical expertise helped keep the project on a sound technical footing. He was always available to answer my technical questions. His reviews of the original proposal and manuscripts greatly raised the quality of this Short Cut.

I would like to give special thanks to my editor, Brett Bartow. He patiently worked with me through the process in generating the proposal, arranging reviews, and keeping things in order and on schedule.

I want to thank all the reviewers who assisted in various stages of the project; their efforts resulted in a much improved final product. Rudy Davis and Tim Gimmel reviewed many chapters under a tight timeline. Tony Phelps, Faraz Shamim, John Skochenski, and Catherine Yan reviewed and provided comments on one or more chapters. Sam Chang, David Friedman, Tony Phelps, and John Skochenski reviewed and provided suggestions for the original proposal.

**FIGURE 3-3**
Simple example of ONS proxy ARP service

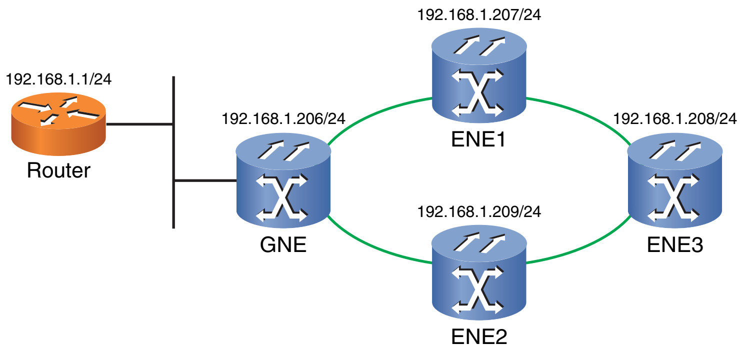


Successful ping responses to both ENEs from the router result in the ARP cache, as shown in Listing 3-1. The GNE's MAC address is mapped to all three IP addresses, an indication that GNE is performing PAS for the three ENEs. In other words, all three ENEs are PAS clients.

**LISTING 3-1 Router's ARP Table**

```
Internet  192.168.1.206        97   0010.cfd2.2ed2  ARPA   FastEthernet1
Internet  192.168.1.207       220   0010.cfd2.2ed2  ARPA   FastEthernet1
Internet  192.168.1.208       220   0010.cfd2.2ed2  ARPA   FastEthernet1
Internet  192.168.1.209       220   0010.cfd2.2ed2  ARPA   FastEthernet1
```

Before discussing the rules for ONS PAS, it is important to introduce a few more terms. The PAS on a GNE uses the following bits of information:

- OSPF database
- LAN interface status
- Routing table

**NOTE**
The client table is a subset of the peer table. PAS is performed only for clients in the ready state.

The OSPF database is used to build a table of peers that are on the same IP subnet and in the same OSPF area. Peers must be reachable over DCC. If a peer has a LAN interface active, that peer is also called a LAN peer. The table of peers is checked against the routing table. If a peer is reachable via a host route, the peer becomes a client. The result is a client table. The GNE then probes the local LAN for other PAS. If none is found, the GNE becomes the PAS. Then the GNE probes for all clients in its local LAN. The GNE probes the LAN by broadcasting an ARP Request for each client. If there is a response, the GNE assumes that the client is on the LAN and can answer ARP Requests by itself. The client is removed from the client list. If there is no response, the client is moved to the ready state. The GNE sends a gratuitous ARP to the local LAN to update all connected IP devices with the client IP addresses mapped to the GNE's own MAC address.

When there is more than one GNE on the same LAN segment, only one GNE becomes the PAS, and all other GNEs become PAS candidates (some variations are discussed later in the chapter). The selection is based on the MAC address: the GNE that has the highest MAC address is the PAS.

If there are link changes on the DCC side or LAN interface status of other NEs, such changes are communicated via OSPF database updates. PAS tables are affected accordingly. It is important to note here that PAS affects only the traffic inbound from the LAN interface of a GNE to ENEs.

To better understand the inner workings of PAS, the previous discussion is extracted into the following basic rules guiding the PAS operation on ONS. Understanding these rules is important in designing and troubleshooting PAS-related networks, because PAS information is currently not available from any of the management tools except the shell interface.

## Rule 1: A PAS Responds Only to ARP Requests from the LAN for Peer Addresses on the Same Subnet as the GNE

This rule is quite obvious, because ARP is defined per the subnet concept. Because the GNE performing PAS has the subnet information, it can easily check whether an address is on the same subnet. If the address is on the same subnet and is a client in the GNE's client list with a ready state, the GNE responds with its own MAC address. Otherwise, ARP Requests for these addresses are ignored.

## Rule 2: Peers Must Be Reachable by the PAS GNE via a Host Route

All the peers must be in the routing table via a host route. The route can be a route either via a DCC interface or via a static route. In earlier software releases (releases prior to 4.6), PAS supports only the DCC-learned host routes and not the static host routes.

## Rule 3: Peers Must Be on the Same OSPF Area as the PAS GNE, or Share One Common Area if the PAS GNE Is on More Than One Area

If a peer is not in the same OSPF area as the PAS GNE, or both do not share a common area, the peer does not make it into the client table. The GNE does not perform PAS for the peer. The same area requirement is due to the fact that only NEs in the same area exchange the topology information, which includes link protection mechanisms (Chapter 5, "Dynamic Routing," provides more details on OSPF database and link topology descriptions).

Figure 3-4 shows an example where an ENE is on the border of two OSPF areas. Two areas are supported on DCC: area 0 and area 1. The GNE performing PAS does not respond to ARP Requests for ENE2, even though ENE2 is in the same subnet as the GNE. Note that this network configuration by itself is not a problem unless the GNE is required to perform PAS for ENE2.

There are several solutions to the PAS problem here:

- Put all ENEs from the same subnet in the same OSPF area.
- Make ENEs on area 1 a different subnet from the GNE subnet. Appropriate routing entries are required to make the end-to-end connectivity. Chapter 4, "Static Routing," and Chapter 5 provide more details on how to solve routing issues.

- Make ENE1 a GNE, and connect the two LAN segments. This solution will work only if the new GNE (ENE1) becomes the PAS. Refer to Rule 6 for PAS election.

**FIGURE 3-4**
PAS across OSPF areas

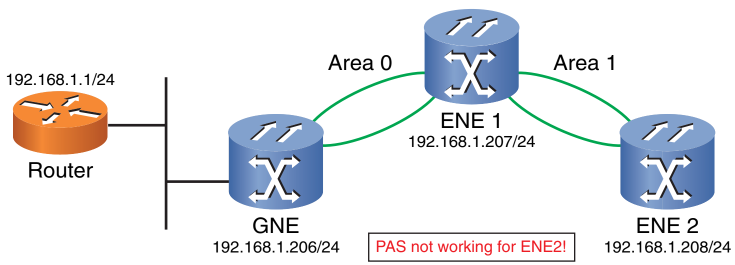


As discussed in Chapter 2, subnetting across OSPF areas is not a recommended practice. Chapter 5 will reexamine the subject from the point of area design.

## Rule 4: A PAS GNE Does Not Perform PAS for LAN Peers Attached to the Local LAN of the GNE

These LAN peers (GNEs on the same LAN, in this case) do not need PAS, because they respond themselves directly to ARP Requests.

## Rule 5: Fragmented LAN Peers May Become PAS Clients

Figure 3-5 shows an example that is the same as that shown in Figure 2-2, but now from the perspective of PAS. Each GNE becomes PAS for its respective LAN segment for other peers. Both GNE1 and GNE2 consider each other a LAN peer (a peer that has a LAN interface active). GNE1 performs PAS for GNE2, ENE1, and ENE2; GNE2 performs PAS for GNE1, ENE1, and ENE2. They do so because they do not discover the other GNE on the same LAN segment. Even though PAS is working fine with the fragmentation for both GNEs, it is not a recommended design because the forwarding is still broken, as described in Chapter 2.

**FIGURE 3-5**
ONS PAS with fragmented subnet

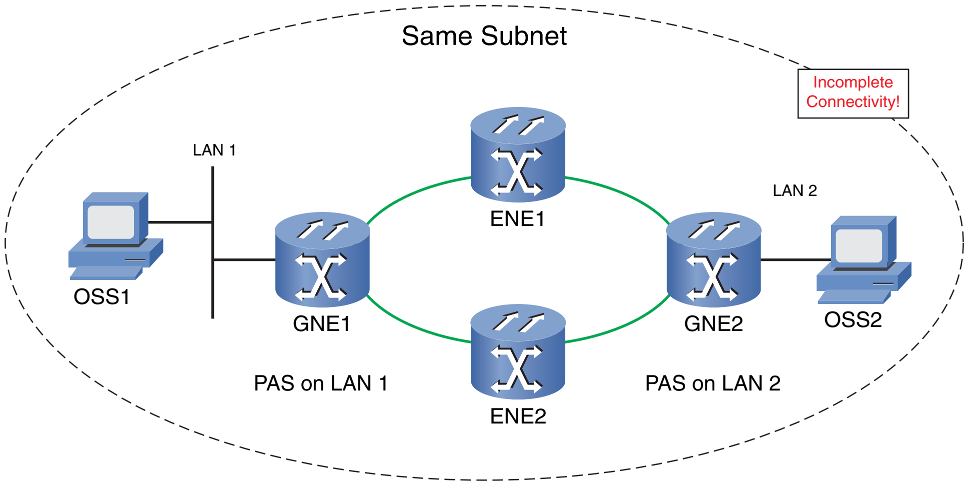


## Rule 6: On a Given LAN Segment, All GNEs Sharing the Same OSPF Database Compete for PAS. Only the One with the Highest MAC Address Becomes the PAS, and All Other GNEs Become PAS Candidates

Figure 3-6 shows an example of this rule. Both GNE1 and GNE2 are on the same LAN segment. They share the same OSPF database because they are on the same OSPF area over the DCC links.

GNE1's MAC address is 0010.cfd2.2ed2, which is greater than the MAC address of GNE2 (0010.cf7d.7e02). GNE1 becomes the PAS for ENE1 and ENE2.

**FIGURE 3-6**
ONS PAS election

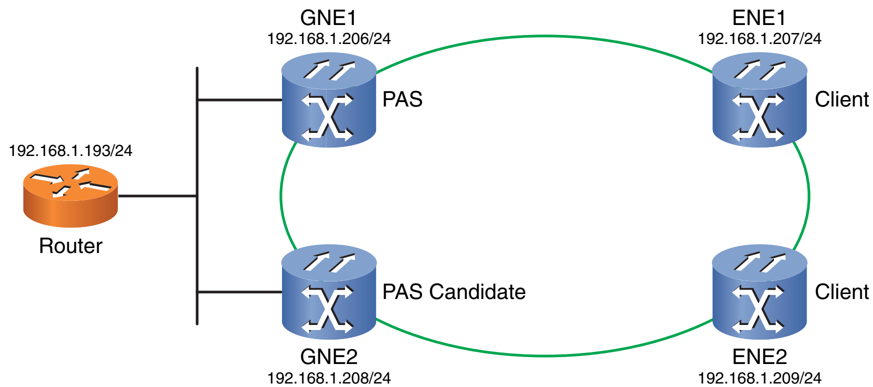


If both GNEs lose visibility over DCC to each other, either by dual-fiber breaks or disabling of OSPF on links connecting the two GNEs, a special case is formed where, although both GNEs are still on the same LAN segment, they do not share the same OSPF database. As a result, the two GNEs do not see each other and the respective ENEs as peers. A new PAS-client relationship is formed.

Figure 3-7 shows the new topology from the point of PAS.

This topology can be achieved by disabling OSPF on either or both DCC ports, pulling the fibers, or moving the IO ports out of service. As a result, GNE1 has only ENE1 as its peer, and GNE2 has only ENE2 as its peer. Because GNE1 and GNE2 do not see each other as peers, there will be no PAS election. Both GNE1 and GNE2 independently become PAS: GNE1 has a client of ENE1, and GNE2 has a client of ENE2. Traffic from Router reaches the intended ENE correctly.

If the topology shown in Figure 3-7 is a result of link failures, the normal PAS operation should return when fiber links are restored. The topology can also be the result of design, where the GNE-ENE link is a linear connection. Regardless, there is no PAS election between the two GNEs.