In this chapter, you learn about the following topics:

- Technologies and Protocols for DSL, Cable, and Ethernet Broadband Networks
- Bridged and PPP Access Mechanisms, with an Evaluation of how well They Solve the Requirements for Broadband, Such as Quality of service, Address Assignment, Service Selection, and so on

CHAPTER 2

Delivering Broadband Access Today: An Access Technologies Primer

A VPN is a service that can carry pure data or multiservice traffic. When you design or implement a VPN for broadband access, you need to understand how the different access architectures can impact design decisions and can actually have some interesting repercussions on the VPN service itself, because of quality of service (QoS) or security trade-offs, to name but two examples.

This chapter reviews the two principal Layer 2 access architectures in use today: bridging and PPP. It looks at how each is implemented on different Layer 1 broadband media, such as digital subscriber line (DSL) and cable. Then it describes how each architecture solves some of the basic requirements of a network service, such as security, QoS support, routing, and address assignment. There are lots of permutations to go through: For example, security on a bridged cable broadband network is different from security on a bridged Ethernet broadband network, so it is worthwhile to look at each case.

If you sometimes feel lost going through all of these different scenarios, remind yourself that you are looking at how common problems are solved on different types of broadband networks. The set of problems is important, because you will want to make sure that a broadband VPN service solves it too. The major topics that are covered are as follows:

- Bridged access architectures
    - — Bridging on DSL using Routed Bridge Encapsulation (RBE), including setup, routing, and address assignment
    - — Bridging on cable
    - — Bridging on Ethernet
    - — Security for bridged access, with a look at different scenarios for DSL, cable, and Ethernet
    - — Authentication and accounting for bridged access
- PPP access architectures
    - — PPP over Ethernet (PPPoE), including setup, routing, and address assignment
    - — PPP over ATM (PPPoA)

— PPP address assignment

— PPP authentication, accounting, and security

Bear in mind that this is a review chapter. If you are comfortable with PPP and bridging, then you can safely skip ahead to the next chapter.

Architecture 1: Bridged Access Networks

Bridged access networks are so named because they transport Ethernet frames transparently across a network. Ethernet is the most successful LAN protocol ever. It has basically replaced all other forms of Layer 2 encapsulation in enterprise networks and is arguably in the process of doing the same thing in residential networks. Not so very long ago, subscribers connected to the Internet directly from their PC using a modem. Today, home networks use Ethernet. For example, laptops have built-in Gigabit Ethernet ports, and wireless LAN is very quickly proving to be an alternative to running cables between rooms the world over. All these scenarios use Ethernet framing, and the most cost-effective broadband service will be Ethernet-centric: Ethernet ports are cheaper, Ethernet cards are cheaper, and Ethernet equipment is cheaper, too. If Ethernet is the user-to-network interface of choice, broadband access networks need some way to carry Ethernet traffic from the subscriber premises to their destination. The easiest way to do this is to simply bridge the traffic—after all, bridging was invented to connect Ethernet LAN segments together, so it should be a pretty useful way to carry Ethernet over WAN connections, too.

However, most broadband networks are not Ethernet based, so the Ethernet frames transmitted by a device on a home network must be converted to some other form before being carried over the native transport medium. For that very reason, bridging in a DSL environment is tricky, because today's DSL uses ATM as the modulation layer of choice. To carry Ethernet, you have to do a form of RFC 2684 bridging. Cable networks are easier because they can natively encapsulate Ethernet frames directly over Data Over Cable Service Interface Specification (DOCSIS). Of course, the simplest scenario of all is one where the access network is Ethernet based, using either standard Ethernet Layer 1 or some form of optical transport over longer distances.

The advantage of bridged architectures is their simplicity. The customer premises equipment (CPE) has no difficult tasks to perform, so it can be very cheap. The overall simplicity has one significant cost, however: namely security. Unless the router enforces some form of security mechanism, all bridged subscribers are in the same broadcast domain and everyone in that domain can see sensitive traffic such as ARP requests, Windows neighbor discovery packets, and so forth. This is a situation best avoided.

Although DSL, cable, and residential Ethernet networks each use radically different transport mechanisms, the issues and design considerations of bridged access are common across all the different media. The next section looks at the details of bridging in DSL networks. This is probably the most complicated scenario, because of the conversion back and forth between ATM cells. Fortunately, bridging over ATM is well standardized in RFC 2684, and the focus of the following discussion is on how bridging over ATM works on Cisco aggregation routers.

Bridging in DSL Using RFC 2684

RBE is a Cisco implementation of bridged Ethernet over ATM with a separate broadcast domain for every ATM circuit. The CPE is a simple bridge that encapsulates Ethernet frames into ATM cells using the RFC 2684 bridging standard. Figure 2-1 shows a typical RBE architecture.

Figure 2-1 *RBE Architecture*

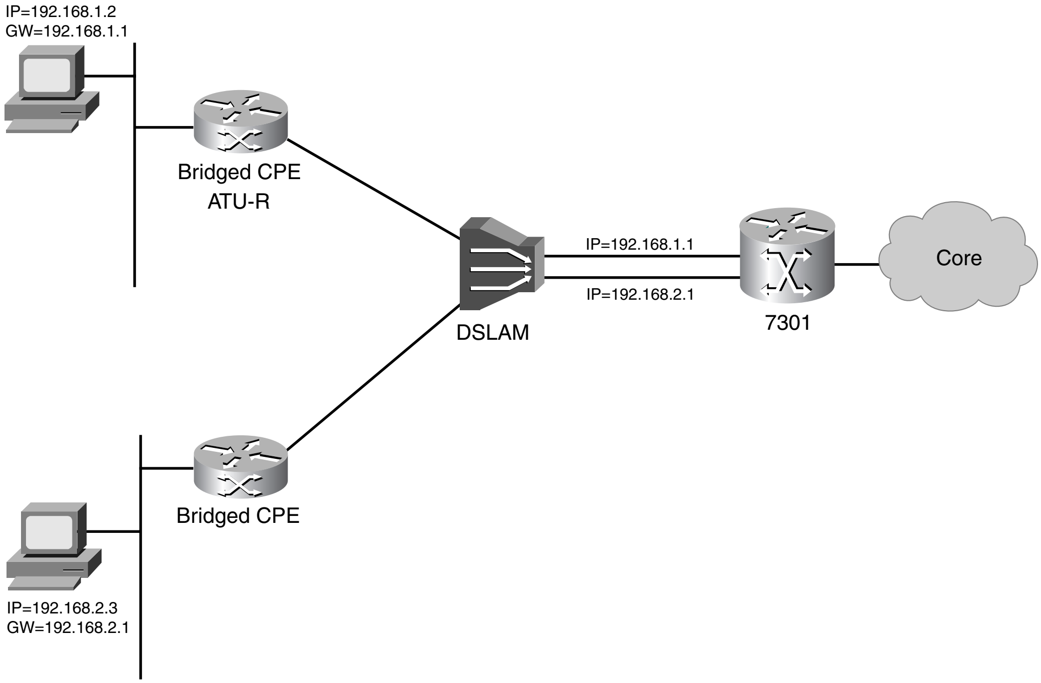

Figure 2-2 illustrates the packet encapsulations used at different points in the network.

Figure 2-2 *RBE Network Cross Section*

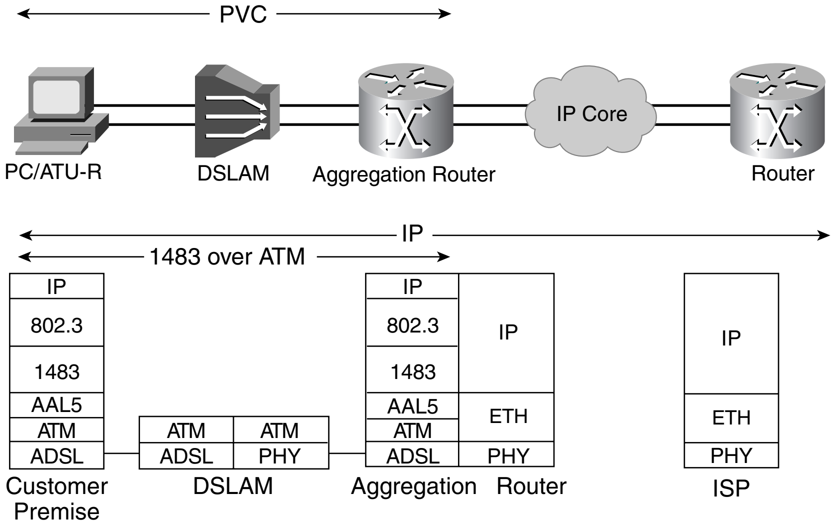

The flow of packets in Figure 2-2 works as follows:

For upstream traffic:

1. The subscriber PC is configured with the aggregation router's IP address as the default gateway. Just as on any Ethernet network, the PC sends an ARP request for the router MAC address and, once it learns it, transmits the Ethernet frame.
2. The bridged CPE encapsulates the Ethernet frame in an AAL5 bridge protocol data unit (BPDU), then segments the BPDU into ATM cells and sends it across the DSL network. Figure 2-3 shows the protocol encapsulations used at different points of the network.
3. The router reassembles the ATM cells, removes the AAL5 information and the Ethernet frame information, and routes the packet to its destination. Note how the router behaves: this is the behavior you would expect to see on a routed interface.

For downstream traffic:

1. A server sends a packet to the subscriber PC that is routed to the aggregation router.
2. The aggregation router has a static route that identifies the interface to use to reach the subscriber's IP address.
3. If necessary, the aggregator issues an ARP request to discover the subscriber PC MAC address. Then it encapsulates the Ethernet frame in an AAL5 Bridged format BPDU, segments everything into ATM cells, and transmits it.
4. The CPE reassembles the ATM cells into AAL5 PDUs, removes the AAL5 information, and transmits the frame on its Ethernet port.
5. The PC receives the data.

Figure 2-3 *Payload Format for Bridged Ethernet/802.3 PDUs (source: RFC 2684)*

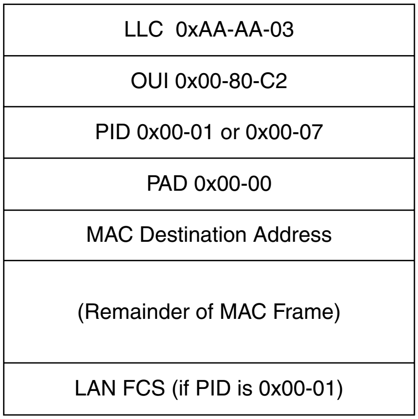

LLC 0xAA-AA-03

OUI 0x00-80-C2

PID 0x00-01 or 0x00-07

PAD 0x00-00

MAC Destination Address

(Remainder of MAC Frame)

LAN FCS (if PID is 0x00-01)

For neighbor-to-neighbor traffic, note that if a subscriber PC sends a packet to another subscriber connected to the same aggregation router, the flow of packets is identical to the upstream and downstream flows described here. It is important to understand that there is no direct Layer 2 path between subscribers and that all traffic must be routed, even when subscribers' circuits are terminated on the same physical port on the router.

In Cisco IOS Software terms, the router in Figure 2-2 uses a logical point-to-point subinterface for each subscriber and treats each of these interfaces as a separate IP network. The default requirement of such a topology is, of course, to have a different IP subnet on every link. But in broadband, you have to manage very large numbers of connections, and there can be thousands of RBE subscribers connected to a single router. In such a case, IP addresses can run out very quickly. To get around this, you use unnumbered interfaces.

When using unnumbered interfaces, as in Frame Relay networks, the router no longer can know which interface to use to send traffic to a particular subscriber just by looking at the destination IP address, because no IP address space is associated with a subinterface. Additionally, to save IP address space, the subscriber IP addresses belong to the same subnet. Therefore, there must be an explicit route statement that maps the subscriber virtual circuit to its IP address. This is why Step 2 for downstream traffic mentions a route—because use of the unnumbered link.

Before learning about RBE configuration, you should understand the alternative to RBE, called *integrated routing and bridging (IRB)*, because both RBE and IRB are used (although there is less and less use of IRB). IRB is a multipoint topology in which all the subscribers are terminated on a point-to-multipoint interface. The architectural problem with IRB is that all of the subscribers are on the same Layer 2 network and are thus part of the same broadcast domain, which makes a network open both to performance degradation because of broadcast storms and to security issues. RBE is a superior, more secure implementation. For example, ARP spoofing is not possible with RBE because an ARP request for a particular address is sent only on the subinterface for that address. With IRB, the request would be flooded to all interfaces in the bridge group. RBE also prevents MAC address spoofing, again because there is a distinct subnet for each subinterface. If a hostile user tries to hijack someone else's address by injecting a gratuitous ARP packet (using their MAC and the victim's IP address), Cisco IOS will detect a subnet mismatch and generate a "Wrong Cable" error.

Note that, from a subscriber's point of view, they both look exactly the same.

Now that you understand the theory and architecture, you are ready to look at some configuration scenarios:

- Basic RBE configuration
- RBE QoS
- RBE routing
- RBE IP address assignment

These sections all get into the details of Cisco IOS commands.

RBE Configuration

RBE router configuration in Example 2-1 is straightforward.

Example 2-1 *Basic RBE Configuration*

```
interface Loopback0
 ip address 192.168.1.1 255.255.255.0
 no ip directed-broadcast
!
interface ATM0/0/0.132 point-to-point
 ip unnumbered Loopback0
 no ip directed-broadcast
 atm route-bridged ip
 pvc 1/32
  encapsulation aal5snap
!
interface ATM0/0/0.133 point-to-point
 ip unnumbered Loopback0
 no ip directed-broadcast
 atm route-bridged ip
 pvc 1/33
  encapsulation aal5snap
```

The configuration is very similar to the regular IP over ATM on a Cisco router, with only the addition of **atm route-bridge ip** to enable RBE. The subscribers' hosts must be configured to use the aggregator interface as their default gateway, in this case 192.168.1.1.

RBE Quality of Service

RBE has a full range of QoS options. Because it runs over an ATM PVC, you can fully exploit all the capabilities of the ATM layer to offer different QoS profiles to subscribers. You need to remember that ATM class of service (CoS) is applied to any and all traffic on the circuit. You can't restrict it to an individual application or destination.

Additionally, you can also use IP QoS. Cisco IOS has numerous bells and whistles that let you apply policies to combinations of application flows, IP destinations, and so forth. You can classify packets, police their rate, queue them, prioritize them—whatever it is you need to do to have different levels of service made available to user applications. You enable IP QoS by applying a **service policy** to a PVC. Example 2-2 adds a **PREMIUM** policy to output IP traffic on PVC 1/33. The **PREMIUM** policy is not included here, but is defined using standard Cisco IOS Modular QoS CLI (MQC) syntax.

Example 2-2 *RBE with IP QoS*

```
interface ATM0/0/0.133 point-to-point
 ip unnumbered Loopback0
 no ip directed-broadcast
 atm route-bridged ip
 pvc 1/33
  encapsulation aal5snap
  service-policy output PREMIUM
```

RBE supports Weighted Random Early Detection (WRED), low-latency queuing (LLQ), and policing.

RBE Routing

As previously mentioned, RBE uses unnumbered point-to-point subinterfaces with a route to each subscriber IP device. Example 2-3 shows the Cisco IOS routing commands, with the required static route for each subscriber.

Example 2-3 *RBE Static Routes*

```
! network routes
router ospf 100
redistribute static
192.168.14.0 0.0.0.255 area 0

!subscriber routes
ip route 192.168.1.2 255.255.255.255 ATM0/0/0.132
ip route 192.168.1.3 255.255.255.255 ATM0/0/0.133
```

Example 2-3 has just three lines of subscriber static routes, but imagine a configuration with 10,000 subscribers connected to the same router. If you announce all these host routes across an IP backbone, you can run into trouble with route table sizes, because each individual entry in a route table consumes memory.

Figure 2-4 shows a simple network with RBE subscribers connected to an aggregation router. Consider the flow when the subscriber PC in this network pings the server at 198.133.219.25.

Figure 2-4 *RBE Packet Flow*

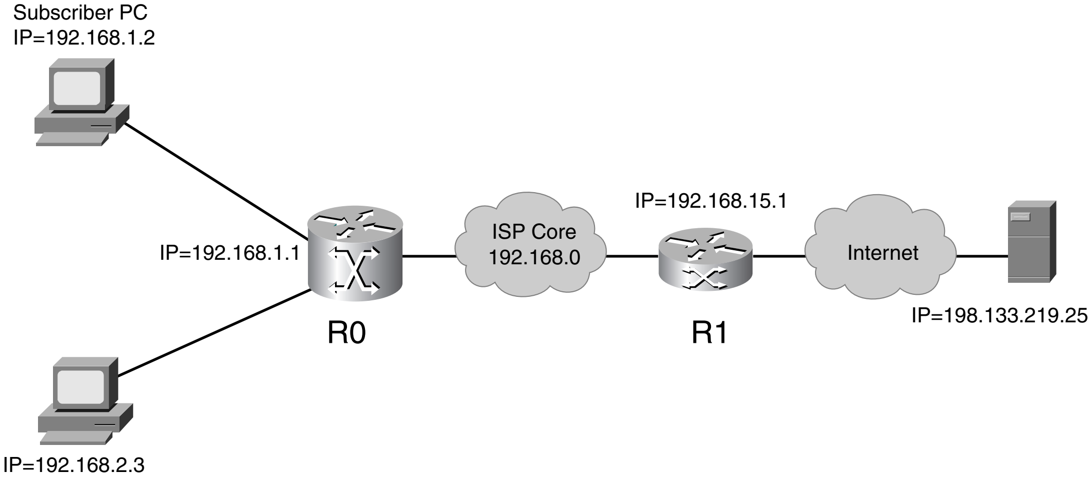

In Figure 2-4, the following happens:

1. The subscriber PC uses the aggregation router's address as the default gateway, so the ICMP ECHO request packet is sent to R0.
2. The aggregation router, R0, also has a simple default route that points to R1. R0 forwards the ICMP packet received from 192.168.1.2 to 192.168.15.1, which is the address of R1.
3. R1 forwards the packet out its egress interface across the Internet. Assuming that IP routing is functioning correctly, the ICMP packet will eventually reach its destination at 198.133.219.25.
4. Now, consider how the ICMP reply is routed back to 192.168.1.2. R1 has announced the 192.168.0.0/16 network to the Internet, so data sent to 192.168.1.2 will reach it.
5. R0 has announced /32 routes for all the RBE subscribers, so R1 will find a route entry for 192.168.1.2/32. (Of course, the next hop will be some intermediary router between R1 and R0.)
6. R1 sends the ICMP packet, which is routed to R0. R0 now looks up 192.168.1.2 in its routing table and finds the static route that points to PVC 1/32 on ATM interface 0/0/0.

Multiply this scenario by several thousand RBE interfaces on several hundred aggregation routers; throughout the 192.168.0.0 network, they quickly grow to unmanageable sizes. It is a much better design to aggregate the routes as soon as possible, and using simple a static route to Null0 is a way to do this. (Plenty of other ways exist, such as configuring OSPF to announce subnet addresses, but traffic to an interface, even Null0, is processed very quickly on a router, so performance is quite good using this method.) In the example shown in Figure 2-4, the ISP 192.168.0.0 backbone routers now need to carry only a single announcement for network 192.168.1.0/24, as indicated in Example 2-4.

Example 2-4 *RBE Static Routes on R0 with* Null0 *Route*

```
! subscriber routes
ip route 192.168.1.2 255.255.255.255 ATM0/0/0.132
ip route 192.168.1.3 255.255.255.255 ATM0/0/0.133
ip route 192.168.1.0 255.255.255.0 Null0
! default route
ip route 0.0.0.0 0.0.0.0 ATM1/0/0.100
```

Now, suppose host 192.168.1.2 sends a ping to 198.133.219.25. The following would happen:

1. On the path from 192.168.1.2 to 198.133.219.25, everything happens as before.
2. On the return path, the server at 198.133.219.25 replies with an ICMP REPLY, which finds its way to R1.
3. R1 has a route for 192.168.1.0/24 that was originally announced by R0. The ICMP REPLY packet will be forwarded to R0.
4. R0 has the same static route to PVC 0/132 on ATM interface 0/0/0.
5. Any traffic received by R0 that is for the 192.168.1.0/24 subnet, for which there is no static RBE route, will be forwarded to Null0 (i.e., dropped).

Although the use of route aggregation is well understood in large IP networks, it has not been widely used in DSL wholesale scenarios, where traffic is tunneled, not routed, to the ISP. Route aggregation is one of the challenges that reappears with IP VPNS and will be discussed further in later chapters.

In the next section, you will see how to create the subscriber routes automatically, instead of statically as in this section. Keep in mind, however, the importance of being able to aggregate as early as possible: You don't want tons of /32 routes wandering around your network.

RBE Address Assignment

You have seen RBE subscribers in all the examples so far with addresses already configured. How did they get them? How do you scale address assignment methods for broadband networks?

Because the preceding sections are all about bridging, DHCP is the logical choice for dynamic address assignment. Statically configuring all the end-station addresses obviously is impossible—the headaches this would create would completely outweigh any employment protection advantages for network operations staff.

When using DHCP on a DSL network, you have two basic options:

- Configure a DHCP server on the aggregation router. This configuration is less common, but entirely possible. You will see configuration examples of Cisco IOS DHCP servers in the “Cable CMTS” section, later in this chapter.
- Use a central DHCP server to which the aggregation router forwards DHCP requests. In this case, the router behaves as a DHCP relay agent.

To do DHCP relay, add the **ip helper-address** command to *every* subscriber interface. The **ip helper-address** command gives the address of the DHCP server, as shown in Example 2-5.

Example 2-5 *RBE Configuration with DHCP Relay*

```
interface Loopback0
 ip address 192.168.1.1 255.255.255.0
 no ip directed-broadcast
!
interface ATM0/0/0.132 point-to-point
 ip unnumbered Loopback0
 ip helper-address 192.168.2.100
 no ip directed-broadcast
 atm route-bridged ip
 pvc 1/32
  encapsulation aal5snap
```

The sequence of events when using DHCP relay is as follows:

1. When the subscriber host starts up, it broadcasts a DHCP Discover packet. This is carried in a BPDU to the aggregation router. The aggregation router recognizes this as the DHCP packet and knows that it needs to forward it to a DHCP server because of the **ip helper-address** on the subinterface. In this case, the aggregator is behaving as a DHCP relay agent.
2. The relay agent actually converts the DHCP broadcast into a unicast IP packet to the DHCP server located at 192.168.2.100. The relay agent puts its own address in the giaddr field of the DHCP packet and puts the subscriber VPI/VCI in the Option 82 field. (This data also includes the receiving interface name, so it is unique per device. The combination of the giaddr IP address and Option 82 yields a globally unique circuit ID.) You can use the global **rbe nasip** command to set the interface address the router puts in the giaddr field.
3. You can configure multiple DHCP servers by entering additional **ip helper-address** commands.
4. The DHCP server returns a DHCP Offer packet.

5 The PC chooses from the different servers that replied to its Discover message and sends a DHCP request to one of them. Remember, the PC still does not have an IP address at this point, so it broadcasts.

6 The DHCP relay agent again forwards the packet to the DHCP server. The relay agent does have an IP address, so it unicasts the packet to the server using UDP.

7 The DHCP server selects an address from an appropriate pool of IP addresses (it can use either the requesting MAC address or giaddr and Option 82 to select the scope) and returns a DHCP reply to the PC, which now has its own IP address and with it the default gateway address.

8 Crucially, the router dynamically creates a host route for the new IP address. As DHCP replies are sent by the DHCP server, the aggregation router looks at the address being assigned and creates a host route to that address using the interface on which the request was originally received. This is one of the bits of magic needed for large-scale deployment. These routes are marked as **static** in the output of the **show route** command.

You should still use the summarization technique discussed with dynamic addresses also. Announce in your favorite dynamic routing protocol the subnet of addresses that you know will be allocated to DHCP requests originating from a particular aggregation router. As new hosts connect to the network, host routes for them are created automatically as soon as they are assigned addresses. The aggregator will then have an aggregate route to make sure that packets are sent to it for the group of potential subscribers, and specific host routes for hosts that are actually active. This way you have the best of both worlds: An aggregate route is announced to peer routers, but per-subscriber routes are dynamically created as IP addresses are assigned.

More Bridged Access—Cable and DOCSIS

The worlds of cable and DSL have some major differences, but from an IP perspective they are very similar. If you ignore the many Layer 1 details on a cable headend router, the router configuration is similar to RBE, and thus many of the points already introduced for RBE also apply to cable access.

Cable modems communicate with headend routers, called CMTS, over the HFC plant using the DOCSIS standard:

> *Data is modulated and demodulated using the North American DOCSIS specifications, with downstream 6-MHz channels in the 54- to 860-MHz range and upstream ranges of 5 to 42MHz. The cable interface supports NTSC channel operation, using standard (STD), Harmonic Related Carrier (HRC), or Incremental Related Carrier (IRC) frequency plans conforming to EIA-S542.*
>
> *NTSC uses a 6MHz-wide modulated signal with an interlaced format of 25 frames per second and 525 lines per frame. NTSC is compatible with the Consultive Committee for*

> *International Radio (CCIR) Standard M.PAL, used in West Germany, England, Holland, Australia, and several other countries.*
>
> *The DOCSIS radio frequency (RF) specification defines the RF communication paths between the CMTS and CMs (or CMs in STBs). The DOCSIS RF specification defines the physical, link, and network layer aspects of the communication interfaces. It includes specifications for power level, frequency, modulation, coding, multiplexing, and contention control.*

This DOCSIS standard is extremely rich, but, at a very high level, provides a TDM-like system of time slots on a shared infrastructure. In the downstream direction, variable-length MPEG-4 frames carry Ethernet frames. In the upstream direction, fixed-length time slots are assigned to the cable modems by the CMTS. The downstream bandwidth is policed by the CMTS according to the QoS profile of each subscriber. Standard Ethernet 802 LLC is run on top of the DOCSIS layer. Figure 2-5 shows the encapsulation stack.

Figure 2-5 *DOCSIS Protocol Stack*

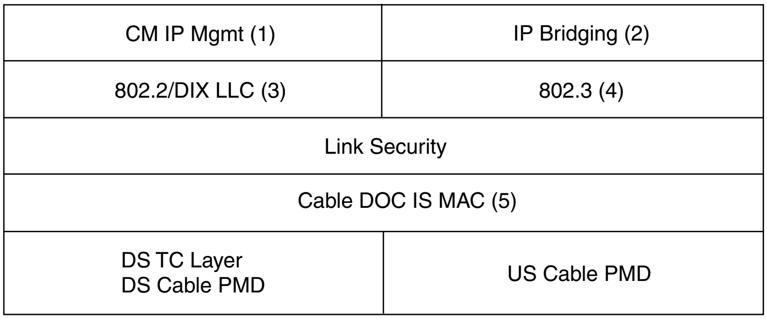

1. DHCP, DNS, TFTP, TOD, CM Registration
2. IP QoS (ToS), IP Classifiers
3. MAC Management Frames, including SYNC, UCD, MAP, …
4. MAC Packet Data Frames
5. SID, …

As part of the session negotiation process, a Service Identifier, or SID, which is part of the DOCSIS Cable MAC layer, is allocated to each cable modem (CM). This is used somewhat like the ATM Circuit Identifier in DSL networks. All traffic sent to and by a given cable modem uses the same SID. The DOCSIS 1.1 specification enhances this to allow cable modems to use several SIDs, each with a different QoS profile so that voice or video can be run over the same infrastructure as data traffic. Again, the parallel with ATM PVCs is apparent.

The cable modem bridges traffic from its LAN Ethernet port over the WAN DOCSIS interface to the CMTS. Subscriber hosts see a shared-access Ethernet network. For upstream traffic, they behave just as RBE clients and need to ARP for the CMTS MAC address. The CMTS also ARPs for PC MAC addresses in the downstream case. Figure 2-6 shows the encapsulations used at different points in the network.

Figure 2-6 *DOCSIS Network Cross Section*

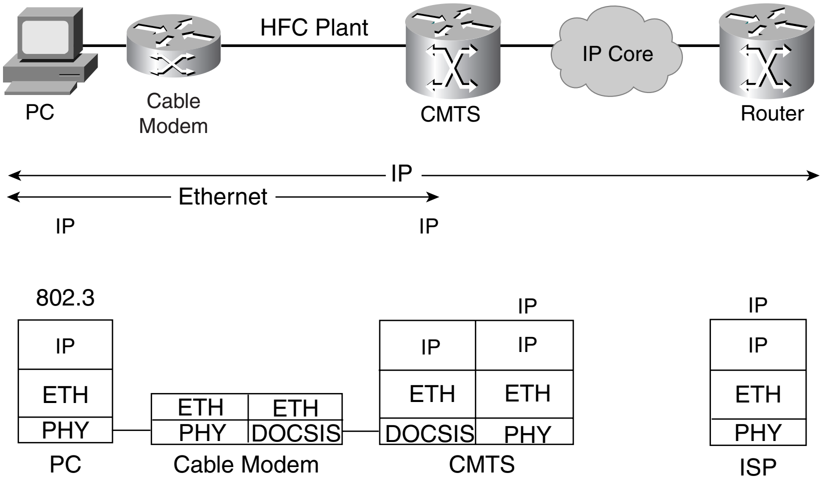

DOCSIS Cisco IOS Configuration

From a Cisco IOS perspective, there are commands specific to the cable plant (HFR) interfaces, cable-modem profiles, etc. Unlike RBE, cable interfaces are natively point to multipoint, which is less secure than point to point. The CMTS and CM have other techniques. Another difference between basic cable and RBE configuration, illustrated in Example 2-6, is the widespread use of secondary addressing (which is also supported with RBE, but is not used very much). In Example 2-6, the primary subnet is for the cable modems; the secondary subnet is for the hosts.

Example 2-6 *Basic Cable Router Interface Configuration*[2]

```
interface Cable4/0
ip address 10.1.1.1 255.255.0.0
ip address 200.1.1.1 255.255.0.0 secondary
load-interval 30
no ip directed-broadcast
cable helper-address 200.1.162.170
no keepalive
cable downstream annex B
cable downstream modulation 64qam
cable downstream interleave-depth 32
 cable downstream frequency 525000000
Cable upstream 0 power-level 0
no cable upstream 0 shutdown
Cable upstream 0 frequency 37008000
cable upstream 1 shutdown
cable upstream 2 shutdown
cable upstream 3 shutdown
cable upstream 4 shutdown
cable upstream 5 shutdown
```

Cable-modem profiles are an important component of DOCSIS networks. These profiles contain configuration instructions, such as upstream and downstream bandwidth, the number of allowed hosts per connection, etc. Example 2-7 shows four different profiles.

Example 2-7 *Cable-Modem Profiles*[3]

```
!
cable config-file platinum.cm
   service-class 1 max-upstream 128
   service-class 1 guaranteed-upstream 10
   service-class 1 max-downstream 10000
   service-class 1 max-burst 1600
   cpe max 10
   timestamp
!
cable config-file gold.cm
   service-class 1 max-upstream 64
   service-class 1 max-downstream 5000
   service-class 1 max-burst 1600
   cpe max 3
   timestamp
!
cable config-file silver.cm
   service-class 1 max-upstream 64
   service-class 1 max-downstream 1000
   service-class 1 max-burst 1600
   cpe max 1
   timestamp
!
cable config-file disable.cm
   access-denied
   service-class 1 max-upstream 1
   service-class 1 max-downstream 1
   service-class 1 max-burst 1600
   cpe max 1
   timestamp
```

Cable Address Assignment

Given that cable broadband is a bridged environment, it shouldn't be surprising to learn that DHCP is used for address assignment. There is a small quirk, though. Even though it functions as an Ethernet bridge, the cable modem also needs an IP address so that it can be managed and it can retrieve its configuration profile. Cable standards mandate the use of Trivial File Transfer Protocol (TFTP) and TOD protocols to retrieve configuration files. TFTP needs an IP address, so the modems use DHCP to get one.

Subscriber hosts also use DHCP to get their addresses. However, for security reasons, the end stations are typically on a different IP subnet than the modems.

The DHCP function on the CMTS router is quite sophisticated. You can either configure the CMTS to relay the DHCP requests to different servers, depending on whether the modem or host sends the packet, or configure different DHCP pools on the router itself—one pool for the modems, one pool for the subscribers. You can also use a mix of the two approaches.

Because of this, the Cisco IOS commands on the CMTS are a little different from RBE and you use **cable helper-address** instead of the standard **ip helper-address** command, as demonstrated in Example 2-8.

Example 2-8 *Cable Router with Multiple DHCP Relay*[4]

```
interface Cable3/0
 ip address 2.41.1.1 255.0.0.0
 no ip directed-broadcast
 no keepalive
 cable insertion-interval 500
 cable downstream annex B
 cable downstream modulation 64qam
 cable downstream interleave-depth 32
 cable downstream frequency 128025000
 no cable downstream if-output
 cable upstream 0 frequency 28000000
 cable upstream 0 power-level 0
 no cable upstream 0 fec
 no cable upstream 0 scrambler
 cable upstream 0 data-backoff 5 12
 no cable upstream 0 shutdown
 cable helper-address 1.1.1.1 cable-modem
 cable helper-address 2.2.2.2 host
```

In this example, there are two DHCP servers. The 1.1.1.1 server receives DHCP requests from cable modems. The 2.2.2.2 server receives requests from subscriber hosts. If you don't specify the [**host** | **cable-modem**] parameter, all requests are forwarded to a single server.

The **cable dhcp-giaddr** command is another powerful addition to CMTS. It modifies the giaddr field with different relay addresses, as demonstrated in Example 2-9.

Example 2-9 *Using the* **cable dhcp-giaddr** *Command*

```
interface Cable4/0
ip address 172.16.29.1 255.255.255.224 secondary
   ip address 10.1.4.1 255.255.255.0
   cable dhcp-giaddr policy
```

The **policy** parameter instructs the router to use the DHCP pool that matches the primary address for requests from the cable modems and to use the pool that matches the secondary address for host requests.

Broadband Ethernet—Ethernet to the Home/Business (ETTX)

Ethernet is still the brave new world of broadband access, and many aspects of broadband Ethernet continue to evolve as market demand and technical solutions develop.

Ethernet can be used in many different service types, which can create some confusion. An Ethernet access network may offer Layer 2 and Layer 3 VPN services as well as other IP-based services, such as Internet access. Figure 2-7 shows the hierarchy of Ethernet services.

Figure 2-7 *Ethernet Services*

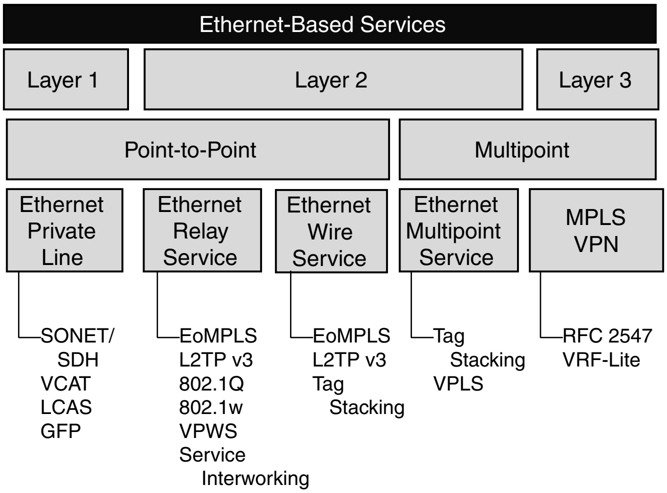

The scenario of interest here, referred to as ETTX, is when Ethernet is used as a last-mile technology for a Layer 3 service, be it Internet or VPN access. The service offering can be for small and medium-sized businesses or residential customers, but today, residential Ethernet is still confined to a metropolitan area.

Unlike DSL and cable, ETTX does not use an existing wiring plant, so it is only cost effective in places where there is an abundance of Category 5 copper cable or fiber, namely multitenant buildings or metropolitan areas. There is a lot of effort today to use Ethernet framing over copper. The next generation of high-speed DSL will, in all likelihood, be Ethernet based. Cisco had an early implementation called Long Reach Ethernet (LRE), covered in more detail in the next section.

Getting back to ETTX, a common residential Internet access architecture, as shown in Figure 2-8, uses a 24-port 10/100 switch as a CPE. Each CPE has a GE trunk port that is ultimately connected to an aggregation router. Figure 2-8 shows a typical configuration that uses switches connected together in a ring between CPE and aggregators that is used to transport Gigabit Ethernet frames.

Figure 2-8 *Residential ETTX*

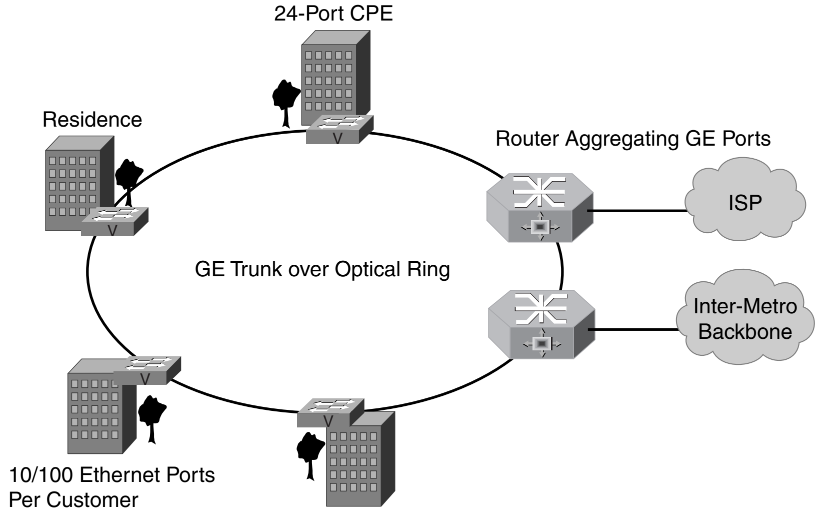
24-Port CPE
Residence
Router Aggregating GE Ports
ISP
GE Trunk over Optical Ring
Inter-Metro Backbone
10/100 Ethernet Ports Per Customer

In a typical network today, there might be 10 CPE on a ring and 20 such rings per aggregator. Although sizable, the number of subscribers per aggregator is still low compared to DSL. Of course, the connection speed is many times higher. The distance from the CPE to the end station is either the standard 100 meters for Category 5 copper cable or several kilometers when using fiber. In the second case, an additional fiber-to-copper converter is required on the customer premises.

The network operation is straightforward and is identical in many ways to a switched Ethernet network in a campus. Figure 2-9 shows the by-now-familiar cross section of protocol encapsulation across the network. Frames are switched from CPE to the aggregation router across intermediary Layer 2 devices. CPE connects to 10/100 Ethernet ports, which are typically trunked over a Gigabit Ethernet port using 802.1q VLANs. The VLANs are terminated on an aggregation router (so they are switched across the access domain), which is responsible for routing traffic to its destination. As with campus networks, each VLAN runs a different IP subnet.

Figure 2-9 *802.1q Network Cross Section*

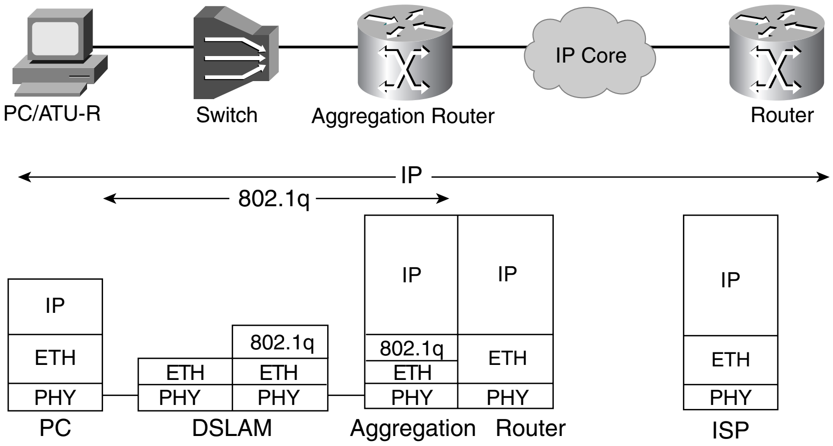

In Figure 2-9, an 802.1q header is inserted in the Ethernet header at the first switch in the access domain. Figure 2-10 shows how Ethernet frames are tagged for VLANs.

Figure 2-10 *802.1q Header*

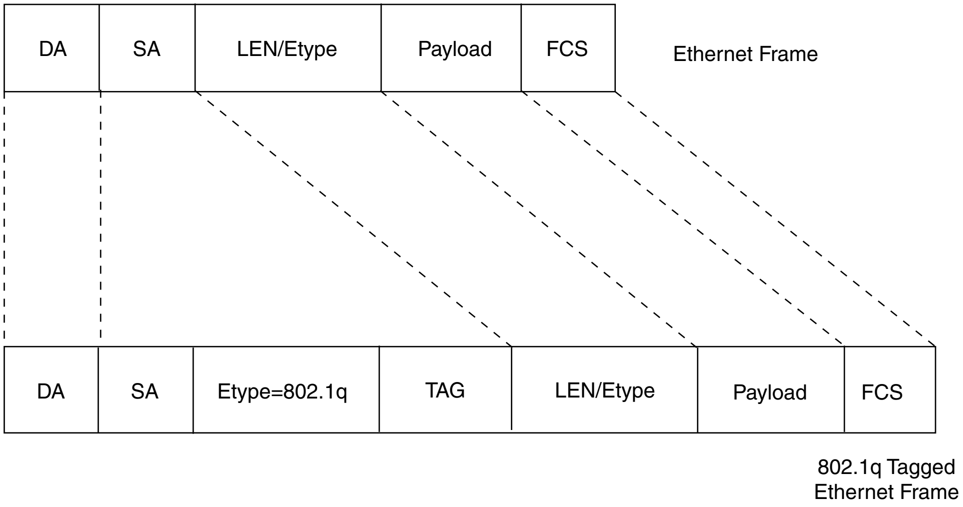

Long Reach Ethernet

Long Reach Ethernet (LRE) is another Ethernet-based solution, this time found within MxUs using copper wiring. Figure 2-11 shows a typical network architecture. Like DSL, LRE can use twisted-pair wiring.

Figure 2-11 *LRE Solution*

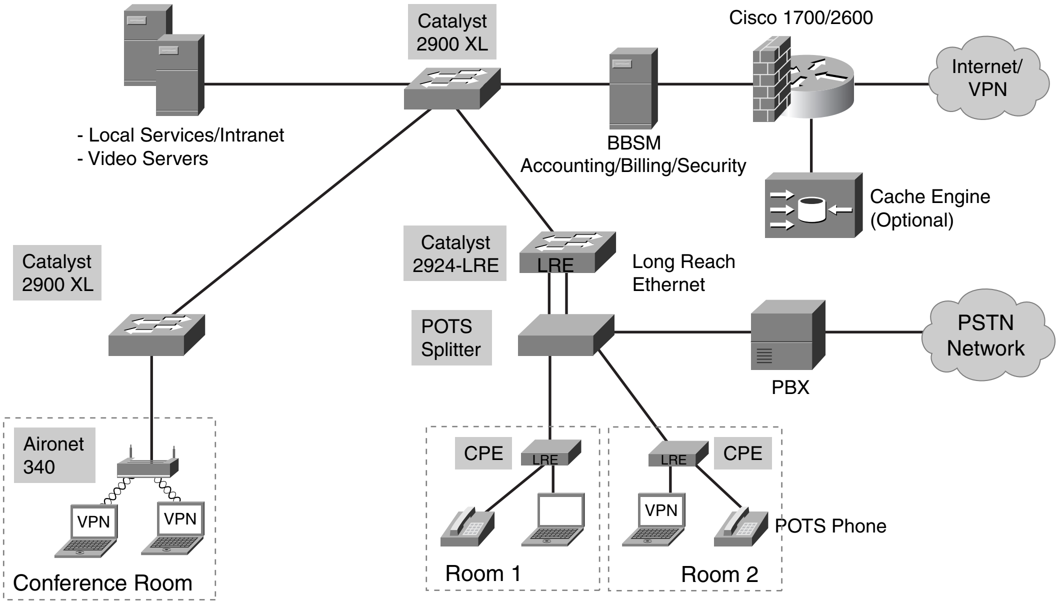

A typical LRE deployment scenario would be a hotel that wants to offer data and voice services to its guests. Because all hotels already have a telephone network, it makes sense to use this expensive infrastructure for data transmission rather than rewire. And because the data traffic is all Ethernet based, it makes no sense to use true DSL, which would require ATM. However, standard Ethernet transmission over telephony-grade wiring is not technically possible, and this is what LRE addresses. LRE uses a transcoding scheme that allows high-speed Ethernet at up to 15 Mbps to be offered across telephony-grade wiring. An LRE CPE encodes the Ethernet frame and transmits it to an LRE-capable switch.

LRE allows voice traffic to either be carried on the same wire using existing analog bandwidth, as with DSL, or to be migrated to IP. Note that LRE works with analog and digital telephones, even if DSL is also used across the same pair of wires. LRE is not widely deployed today as a consumer solution because of ongoing issues with signal interference. It is unclear whether this solution will remain cost effective given the ever-increasing success of wireless Ethernet.

ETTX Configuration

At a basic level, the ETTX configuration is identical to a campus solution. A CPE switch runs 802.1q on a Gigabit Ethernet trunk interface, with each access port in the same VLAN. In the case of Example 2-10, the CPE uses VLAN2.

Example 2-10 *ETTX CPE Configuration*

```
! CPE access port
interface GigabitEthernet0/6
switchport access vlan 2
...

! CPE trunk port
interface GigabitEthernet0/1
switchport trunk encapsulation dot1q
...
```

The Ethernet aggregator configuration has a trunk port and an IP subinterface for every VLAN, as demonstrated in Example 2-11.

Example 2-11 *ETTX Aggregator Configuration*

```
Interface GigabitEthernet1/0/0.22
encapsulation dot1Q 2
ip address 192.168.11.1 255.255.255.240
```

There are networks with a different VLAN for every subscriber, but others in which a different VLAN is used for every service. In the case of per-subscriber VLANs, there must be a scaling mechanism of some kind because the maximum number of VLANs by default is 4000. (The trick is to add a second VLAN tag, known as QinQ.) When there is a VLAN per service, all subscribers are on the same IP subnet.

ETTX Quality of Service

ETTX is often perceived to have the weakest QoS infrastructure of the three access network types under consideration. Although there is no standardized equivalent to ATM's classes (CBR, VBR, etc.) or DOCSIS, Ethernet switches do offer relatively rich QoS capabilities, such as IP- and TCP-based classification, IP DSCP or 802.1p tag-based prioritization, and sophisticated scheduling and policing. Additionally, COS-to-IP DSCP mapping can be done automatically, or COS can be set on a port basis depending on the trust that is ascribed to a user.

Cisco IOS access lists can be used for classification, so different applications or hosts can be treated differently. Even quite low-cost switches can offer multiple queues per port, which is required for multiservice applications.

As anecdotal evidence, remember that a lot of enterprises run voice over their switched infrastructure, which is a testimony to the level of QoS that Ethernet infrastructure can provide.

Apart from transporting multiservice traffic, IP QoS can also be used to help compensate for the fact that Ethernet does not offer many increments for service offerings, with jumps from 10, 100, and 1000 Mbps. Subinterfaces can be policed to lower or intermediate rates, such as 2 Mbps, 34 Mbps, and others, as demonstrated in Example 2-12. You can mark down nonconforming packets, or discard them, to enforce the particular service contract.

Example 2-12 *ETTX QoS Configuration—Policing Subscriber Interfaces*

```
policy-map option-128k
    class class-default
       police 128000 10000 10000 conform-action set-prec-transmit 0 exceed-action drop
policy-map option-512k
    class class-default
       police 512000 10000 10000 conform-action set-prec-transmit 0 exceed-action drop
policy-map option-1Meg
    class class-default
      police 1000000 10000 10000 conform-action set-prec-transmit 0 exceed-action drop
policy-map option-10Meg
    class class-default
        police 10000000 10000 10000 conform-action set-prec-transmit 0 exceed-action
drop
interface GigabitEthernet1/0/0.15
    desc VLAN connecting to customer1
     encapsulation dot1Q 15
     ip address x.x.x.x y.y.y.y
    service-policy input option-128k
    service-policy output option-128k
Interface GigabitEthernet1/0/0.88
    desc VLAN connecting to customer2
     encapsulation dot1Q 88
     ip address a.a.a.a b.b.b.b
    service-policy input option-1Meg
    service-policy output option-1Meg
```

ETTX Address Assignment

Unsurprisingly, ETTX uses DHCP for address assignment to simplify address management and distribution. If, as is common today, the network is owned and operated by a single entity, there are no new issues related to address assignment beyond those already discussed thus far in the chapter. Addresses are assigned using DHCP, and the Ethernet CPE in Figure 2-8 adds option 82 information if port identification is required. The role of the CPE is very important for security, as discussed in the next section.

Using IP addresses efficiently is just as important in ETTX networks as in any DSL or cable network. Consider the case in Figure 2-12 of an Open Access ETTX network in which each subscriber can belong to one of two ISPs, ISP A or ISP B, both of which use DHCP to assign addresses.

Figure 2-12 *Open Access Architecture for Residential Ethernet*

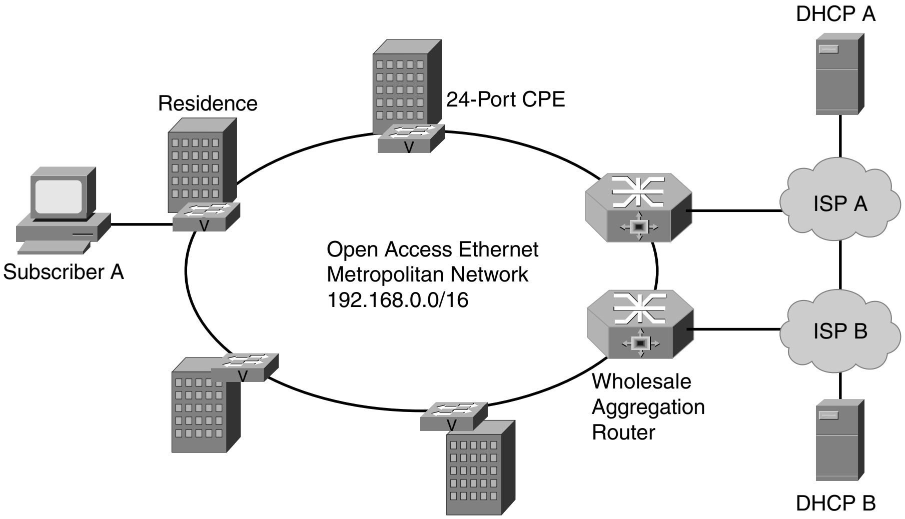

Ethernet can be delivered over point-to-point or ring topologies. Although Figure 2-12 shows just one ring, there can be multiple rings of CPE connected to every aggregation router. Each ring is terminated on a physical interface, with potentially many subinterfaces, each one corresponding to a different VLAN (and there are different policies for how VLANs are used, as previously discussed). As each VLAN corresponds to a different IP subnet, there must be as many subnets as there are VLANs. For an ISP, this can result in wasted address space.

To understand the issue of address waste, consider the following sequence of steps:

1. ISP A has 60 subscribers in the metropolitan region and wishes to use the 192.168.1.0/26 subnet.
2. Subscriber A connects on the first ring and sends a DHCP request, which is relayed to ISP A's DHCP server. VLAN20 is used.
3. The second subscriber connects, but this time in a different part of the city and on a different ring. This time, traffic is in VLAN34.
4. ISP A would like to have the same subnet for all the subscribers in this metro area, but needs a different subnet for VLAN34, or else the aggregation router could not route traffic correctly to subscribers on different subinterfaces.

The bad solution to this problem is to use a different subnet for every VLAN. Unfortunately, ISP A probably has no way of knowing how many subscribers will be on each VLAN and so would potentially need to use as many /26 subnets as there are VLANs in the network. This results in huge waste.

The solution is for the metro service provider to use unnumbered interfaces. Each subscriber has a loopback interface configured for his pool and all the VLANs are unnumbered, as demonstrated in Example 2-13.

Example 2-13 *Unnumbered Interfaces and Loopbacks for ETTX*

```
interface loopback 0
    desc ISP A
    ip address 192.168.1.1 255.255.255.0

interface loopback 1
    desc ISP B
    ip address 192.168.2.1 255.255.255.0

interface GigabitEthernet1/0/0.15
     encapsulation dot1Q 15
     ip address unnumbered loopback 1
     ip address secondary unnumbered loopback 2
Interface GigabitEthernet1/0/0.88
     encapsulation dot1Q 88
     ip address unnumbered loopback 1
     ip address secondary unnumbered loopback 2
```

NOTE Open access is already widely deployed for DSL, but it is still a relatively new concept for ETTX networks. The architecture will continue to evolve. In Example 2-14, the addresses used by the different ISPs can't overlap, because they are terminated on the same router. You will see how to lift this restriction in Chapter 7, "Implementing Network-Based Access VPNs Without MPLS."

Security Considerations for Bridged Broadband Architectures

Security is an important part of an Internet access service, whether it is sold to residential or business customers. Security at the transport layer and application layers is beyond the scope of this work and, indeed, is independent of the type of access used. However, lower-layer security is an important part of overall network design. If the lower layers are not secure, then it is easy for an attacker to work up the stack and compromise application data such as usernames, passwords, credit card numbers, and so on.

The common Layer 2 and Layer 3 risks are as follows:

- **Address spoofing**—This category loosely encompasses all attempts to modify an end station address. It can be something as simple as changing the address on a Linux station NIC or manually changing your IP address.

ARP-based attacks are more sophisticated and, because ARP is not an inherently secure protocol, spoofing is not as difficult as it should be. These attacks involve, for example, sending an ARP reply packet with a spoofed IP address. Most routers or switches simply overwrite an IP address in their ARP table with the IP address obtained from the most recent ARP response, or may record multiple IP addresses from responses, making it easier for the attacker. Slightly more sophisticated is the use of gratuitous ARP, whereby the attacker spontaneously advertises an ARP packet with its MAC address and someone else's IP address. This is completely RFC compliant and all stations that receive the ARP packet happily install the spurious MAC/IP mapping in their ARP tables. Now, no ARP request will be sent when traffic is received for the IP destination. The use of gratuitous ARP is typically used for man-in-the-middle attacks.

As you can gather from the preceding explanation, ARP attacks are limited only to stations that are on the same broadcast domain as the offending attacker.

- **DoS attacks**—DoS attacks can be mounted in many ways. Some simple examples include using gratuitous ARPs to fill the CAM table on an Ethernet switch; using DHCP requests to exhaust IP addresses; or sending a very large number of Layer 3 flows to the default router. DoS attacks often use some form of address spoofing. DoS attacks can be very hard to prevent. Simply constantly changing Ethernet source addresses can be very effective against a switch.

  Sophisticated techniques are under development to improve network security against DoS attacks. Today's routers already have source address checking, access lists, and NetFlow statistics.

- **Broadcast traffic and OS weaknesses**—This is not really a category of network attacks, but more an observation that many host stations are inherently insecure "out of the box" and allow any neighbor machine on the same broadcast domain to browse disk contents. In remote access, the ability to broadcast is trouble.

Security in DSL Broadband Networks

RBE has two characteristics that contribute to network layer security:

- The router uses point-to-point interfaces.
- The broadcast domain is limited to a single site.

If you send a gratuitous ARP packet, the router may install it in its ARP table, but downstream traffic is always directed to the correct PVC because of the host routes used with RBE. When the router receives a packet for a host, it sends the ARP only on the PVC that matches the host route for that address. In other words, ARP is used only after the correct subscriber interface has been identified by the IP routing table. Layer 2 attacks are hard to do successfully in this case.

Even if you do successfully spoof an IP address, the host routes again make sure traffic reaches the correct destination, as described in the following step sequence:

1. Host A sends a packet with a spoofed address of 1.1.1.1 to 192.168.1.100. Its true address is 192.168.1.99.
2. The RBE router forwards the packet to 192.168.1.100.
3. On the return path, the router either routes the traffic to the correct interface for address 1.1.1.1 or drops it if no such address exists in the routing table. It is not returned to host A.

NOTE Before you conclude that this technique is a good way to send large and unsolicited streams of traffic to your neighbor, remember Unicast Reverse Path Forwarding (uRPF) checking. uRPF, if configured, will drop the packet because a router will not accept a packet if the incoming interface is different from the interface defined in the route table that reaches the packet's source address.

Security in Cable Broadband Networks

Because security has long been an issue on cable networks, the Baseline Privacy Interface was added to the DOCSIS specifications to give a secure communication channel between each cable modem and the CMTS.

Unlike RBE, the CMTS router uses a point-to-multipoint interface. Potentially, then, ARP-based attacks are possible because the CMTS sends ARP packets to all hosts on its physical interface. However, remember that the DOCSIS layer also offers protection. The cable modem can store the MAC address/SID mappings, which the network administrator can poll to troubleshoot security issues.

The **cable source-verify** command is important. It configures the CMTS to enforce address assignment by dropping packets with source addresses that it has not seen assigned by a DHCP server (using the **cable source-verify dhcp** option). For this to work properly, the DHCP server must support the DHCPLEASEQUERY message. The **cable source-verify** command prevents attacks based on theft of IP addresses (using a valid address that belongs to someone else) as well as attacks based on invented addresses (either addresses that are valid but not assigned or addresses that are made up).

Security in Ethernet Broadband Networks

Without VLANs, all subscribers in a switched network are in the same broadcast domain, which is an open invitation to trouble, because of the risks that this scenario creates (some of which were discussed earlier in this chapter). Dedicating a VLAN to each customer is, in theory, possible, but it is impractical because of the 4096 global VLAN limit in the 802.1q protocol. It is more common to configure a single VLAN per switch. This still leaves everyone on the same switch in the same broadcast domain. To solve this issue, private VLANs are used.

Private VLANs prevent traffic from a subscriber port from going to any other port on the switch, with the exception of the trunk port, which is defined as a promiscuous port. From the perspective of the subscriber, as represented in Figure 2-13, subscribers "see" a private point-to-point link to the router. The aggregating router, in turn, sees a switched Ethernet segment with multiple subscribers.

Another alternative is to use double VLAN encapsulation, called QinQ, where traffic from each subscriber's port is mapped to a different 802.1q tag, and then that frame is tagged again when it leaves the switch with a tag that uniquely identifies the switch on the Ethernet network. The aggregation router has to be smart enough to handle this double layer of VLAN tags.

Figure 2-13 *Private VLAN*

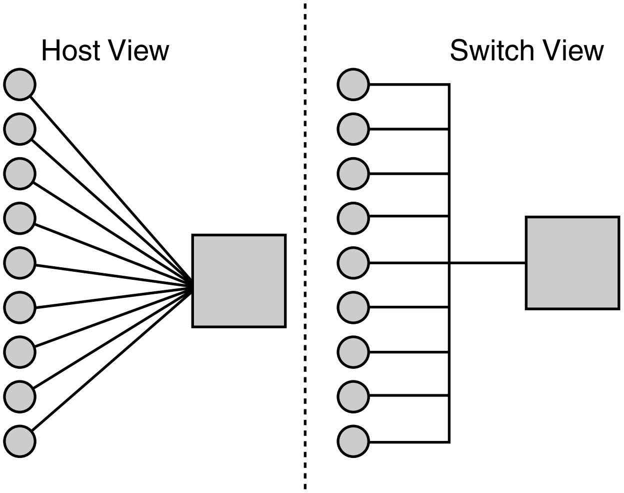

Limiting the broadcast domain stops some of the simplest attacks, but does not prevent ARP or IP address spoofing. Remember that for cable, the router enforces DHCP assignments. If a host tries to change an address, or sends a gratuitous ARP packet, the CMTS ignores it, because it was not assigned to the host by a known DHCP server. The Ethernet scenario is harder to manage because subscriber interfaces have been aggregated on the switch downstream in the network, which may not yet have the necessary mechanisms to enforce Layer 3 to Layer 2 bindings.

Port security is a useful feature that prevents the switch from allowing a MAC address learned on one port to be used on another. Port security can also allow static MAC addresses to be configured for each port. Admittedly, it is hard to scale a solution based on static addresses for a large number of subscribers without an excellent OSS system.

Ethernet switches do have an increasingly large array of tools to deal with DoS attacks. These tools include broadcast suppression, ARP throttling, route processor rate limiting, security ACLs, and so on. On good-quality switches, these features are all implemented in hardware, so they are very efficient and are well worth some extra cost.

NOTE Identification and authorization should always be done as soon as possible in a remote-access network. There is a huge difference in effectiveness between being able to apply policies on a device where each subscriber is on a different port or Layer 2 circuit and one where this is not the case.

Currently, the 802.1x standard is emerging as a possible solution to this problem. 802.1x, which is a port-based access mechanism, works as follows:

1. The client station, or supplicant, sends a request to the switch, or authenticator, which forwards it to a RADIUS authentication server.
2. The authentication server returns a challenge to the supplicant, which must correctly respond to be granted access to the network.
3. The authenticator provides access to the LAN.

This is reminiscent of Point-to-Point Protocol (PPP) authentication and you can think of 802.1x as using a PPP-like authentication mechanism to provide port-based access control.

Authentication and Accounting in Bridged Broadband Architectures

One of the significant differences between Ethernet-based access and the PPP-based solutions is how subscribers are authenticated on the network.

With bridging, there is no authentication using a subscriber's name. This is always ideal, if you have the option, because you authenticate an individual and can then enforce policy with a fine level of granularity.

In bridged architectures, the network-access control is Layer 2 based. If you have a valid MAC address, from a valid port, you will receive a valid IP address. At no time does the user have to enter a name and password. So the user identity must always be tied back to the Layer 3 or Layer 2 addresses and, as you've just seen, these are not the most secure.

Billing is another weakness of the Ethernet solution if the service provider wants to offer a metered service. There is no standards-based way to retrieve usage statistics. Some switches have some useful data in MIBs, but some don't. On routers, if there is hardware for it, you can enable NetFlow accounting. However, NetFlow accounting on high-speed networks generates a considerable amount of data, and the OSS systems must do a lot of data crunching and cross checking between systems to work out which flow belonged to which person at a given time.

Architecture 2: Point-to-Point Protocol Networks

Dial-up access was the first remote-access mechanism and PPP was extensively enhanced to work in this environment. Because of this long deployment experience, PPP is a very mature solution that includes a rich control plane that lends itself to wholesale services. PPP is a more complex protocol than simple bridging, and this additional cost is incurred both on the client stack and on the router.

The following sections look at the variants of PPP used in broadband networks. PPPoE, by far the most prevalent variant today, involves a very simple bridging CPE. As you will see, the PPP session can be initiated from a PC. PPPoA is less common and usually (but not always) the session is initiated from a router.

PPP over Ethernet—The CPE as a Bridge

PPPoE is an interesting protocol. As the name implies, it involves a PPP session running over an Ethernet MAC layer. PPPoE is interesting because PPP was created for point-to-point interfaces, so it needed some enhancements to allow it to run on broadcast media. These enhancements included a discovery process very like the one in DHCP, which serves to establish a logical point-to-point relationship between a PPPoE client and server. All this can sometimes be confusing until you realize that PPPoE is really a superset protocol of PPP in which there is a supplementary setup protocol that runs before regular PPP starts. Cisco IOS debug output shows this really well.

Although the explanation that follows is DSL based, PPPoE is also a perfectly valid possibility for cable and ETTX networks. Figure 2-14 and the list that follows describe PPPoE operation in more detail.

Figure 2-14 *PPPoE Protocol Operation*

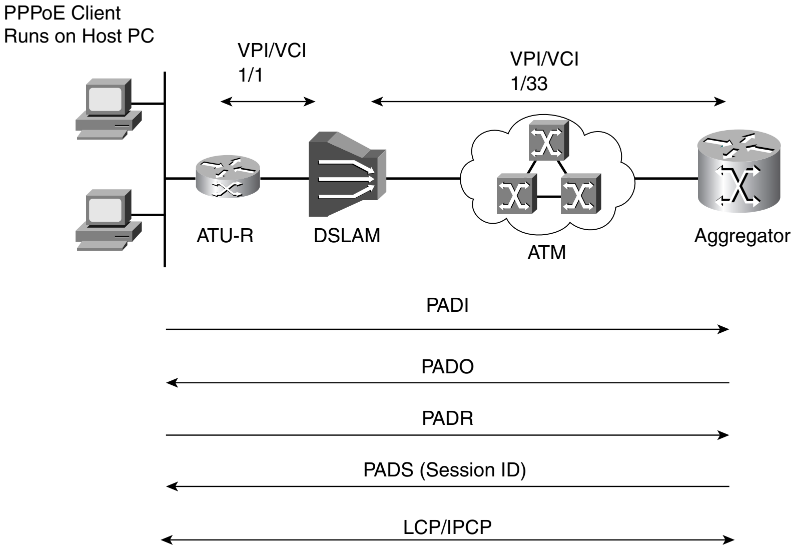

The PPPoE protocol starts before a PPP session comes up, as follows:

1. A host broadcasts an Active Discovery Initiation packet, called a PADI.
2. In theory, a number of different servers (called Access Concentrators in the RFC) can now reply with an Active Discovery Offer packet, or PADO. In the case of DSL, all the Ethernet traffic is bridged across an (point-to-point) ATM virtual circuit by the bridged CPE to a single aggregation router, which is, in fact, the PPP Access Concentrator, so it is hard to imagine a host receiving more than a single PADO on an operational network. However, it is possible. Depending on the implementation of the PPPoE stack on the client, the client may accept only the first PADO returned, or the first PADO returned with the service it wishes to connect to.
3. The host sends an Active Discovery Request (PADR) packet to a single Access Concentrator.
4. The handshake completes when the Access Concentrator sends an Active Discovery Session-confirmation (PADS) packet, which contains a session ID. In the fairly rare case where the Access Concentrator and host are on a broadcast network, this session ID establishes a logical point-to-point connection (The combination of the host MAC address and session ID is unique on the server.)
5. Standard PPP negotiation starts, with LCP and IPCP, just as for point-to-point serial connections.
6. Either side can terminate the session with an Active Discovery Terminate (PADT) packet. This is equivalent to cutting a wire because no more traffic is passed, not even to close the upper-layer PPP session.

Figure 2-15 shows the Ethernet payload for PPPoE.

Figure 2-15 *PPPoE Header (Source: RFC 2516)*

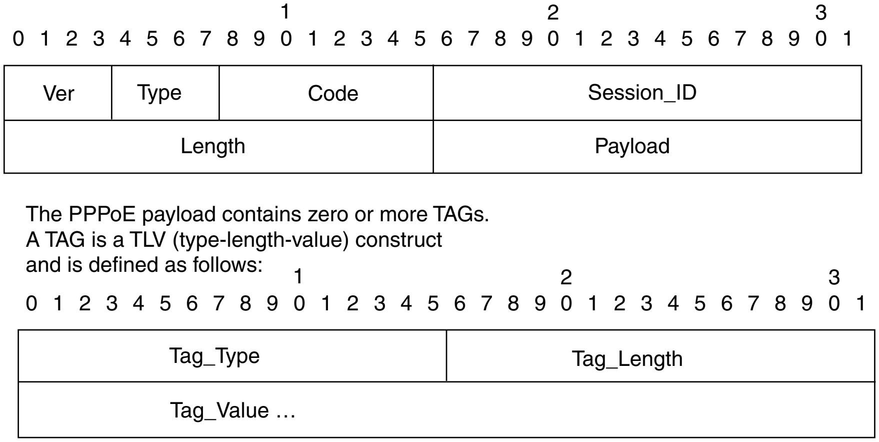

The PPPoE payload contains zero or more TAGs.
A TAG is a TLV (type-length-value) construct
and is defined as follows:

Packet flow is very simple. Each subscriber is terminated on the PPP peer using a virtual-access interface. Just as with dial-up operation in Cisco IOS, these interfaces are cloned from virtual templates. All traffic is routed by the aggregation router to and from each subscriber, often referred to as *hair-pinning*. Host routes are also automatically created for each subscriber. The router treats each virtual access as a directly connected interface, and routing table entries are marked appropriately (with a letter C in a **show ip route** command). Figure 2-16 shows the packet encapsulations used at different points of a PPPoE network.

Figure 2-16 *PPPoE Network Cross Section*

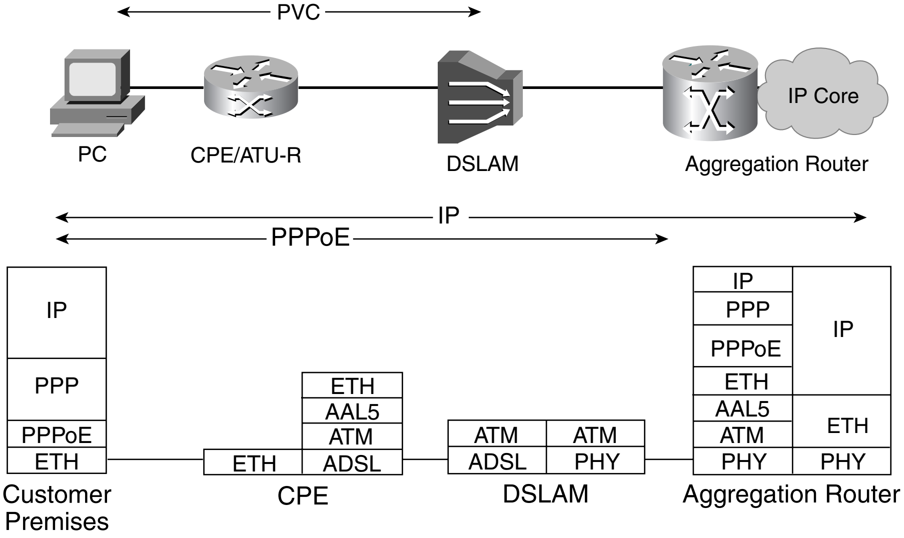

PPPoE Configuration

The PPPoE protocol runs over Ethernet frames. The host sends PPP packets in Ethernet frames, which are in turn segmented by the bridge CPE into ATM cells and sent onward over the DSL network. There are two Ethernet type values for PPPoE: 0x8863 for PPPoE discovery and 0x8864 for the actual PPPoE data sessions.

PPPoE is configured in two parts on a Cisco IOS router. The first part involves creating a PPPoE server. The second part is the standard configuration for PPP. Figure 2-17 shows the typical PPPoE architecture.

Figure 2-17 *PPPoE Architecture*

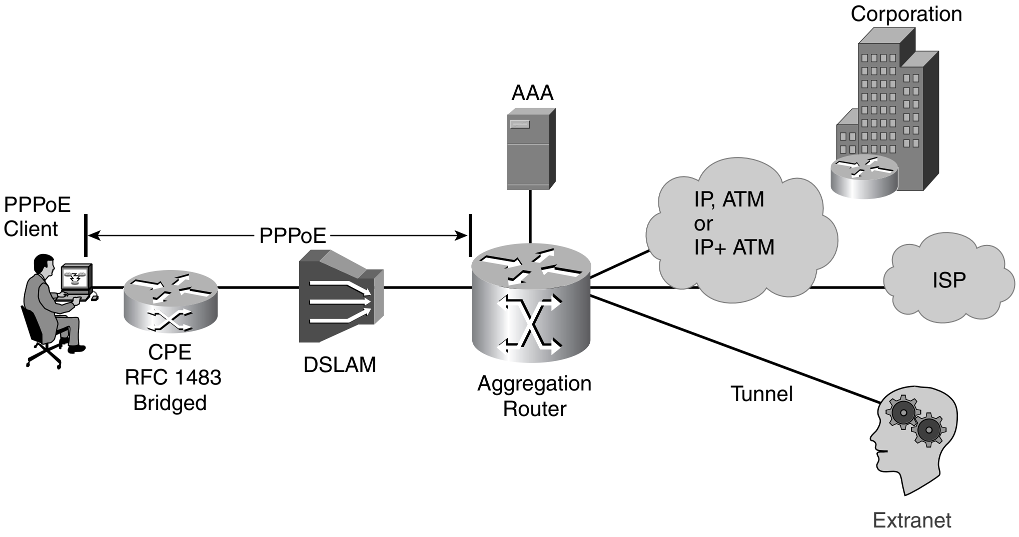

Example 2-14 demonstrates how you would configure a network such as the one shown in Figure 2-17.

Example 2-14 *PPPoE Configuration*

```
vpdn enable
! configuration for pppoe server
! interfaces for sessions will be cloned from virtual-template 1
vpdn-group 1
 accept dialin
  protocol pppoe
  virtual-template 1
! enable pppoe on this subinterface
! subscriber interface
interface ATM0/0/0.132 point-to-point
no ip directed-broadcast
pvc 1/32
 encapsulation aal5snap
 protocol pppoe

! virtual template for pppoe virtual-access interfaces
! note MTU size adjustment
interface virtual-template 1
ip unnumbered loopback0
no ip directed-broadcast
ip mtu 1492
peer default ip address pool pppoe-pool
ppp authentication pap
! pool of addresses for pppoe subscribers
ip local pool pppoe-pool 192.168.10.10 192.168.10.100
```

Example 2-15 shows the PPPoE configuration for the network in Figure 2-16 using Ethernet instead of ATM subscriber interfaces.

Example 2-15 *PPPoE Configuration*

```
vpdn-group pppoe
accept-dialin
protocol pppoe
virtual-template 1
!
interface FastEthernet2/0.2
encapsulation dot1Q 2
pppoe enable
!
interface FastEthernet2/0.3
encapsulation dot1Q 3
pppoe enable
```

NOTE Examples 2-14 and 2-15 use **vpdn-group** commands. In more recent versions of Cisco IOS, this syntax has been changed to use the newer **bba-group** commands. Chapter 6 shows how to use **bba-group**, but a lot of networks still run Cisco IOS images that have **vpdn-group**, which is why those commands are shown here and in Chapter 3, "VPNs in Broadband Networks."

PPPoE Service Selection and Discovery

Another innovation in the PPPoE protocol is the use of PADS messages to advertise services to clients. The premise behind this use is that each host potentially is subscribed to multiple network-based services and needs to choose between them or, alternatively, to discover the list of subscribed or permitted service names. Example services might include a public Internet connection, a private VPN connection, and an extranet managed by a financial institution. To switch between each service, the end customer must have some way of identifying and selecting the service in the first place. Similarly, to subscribe to a brand new service, there has to be some way to inform the customer that it exists.

It is possible to do service selection using PPPoE. To do this, the PPPoE Access Concentrator sends a list of available service names (such as Internet, VPN, SafeShopping) to the PPPoE client. The client then displays the list of services to the user, who can chose whichever one she wants to use.

From an architecture perspective, PPPoE service selection poses an interesting quandary. In the DSL reference model, only ISPs sell services to end users. Yet here is a protocol that only the wholesale providers (the Access Concentrator is a PTT device) can typically use to announce services they do not typically sell to customers they do not typically own.

That said, there is a proposal at the IETF that would allow just the PADS messages to be carried over L2TP. This would allow the ISP's network server to terminate the PPPoE protocol and manage service announcements, which is probably a more logical arrangement from a business perspective. The Cisco implementation is called PPPoE Relay.

PPP over ATM: The CPE as a Router

PPPoA was first standardized at the DSL Forum. It is commonly used, but not as widely as PPPoE. PPPoA is actually simpler than PPPoE because there is no need for any of the extensions, such as discovery.

PPPoA can be run directly from a host with an ATM NIC card. This scenario is fairly rare in operational networks because Ethernet NICs are so much cheaper. Our discussion will focus on a CPE with PPPoA.

To higher-layer protocols, a PPPoA link appears as a routed connection, with remote peer authentication and the possibility of dynamic address assignment.

PPPoA Configuration

PPPoA runs over AAL5 MUX or SNAP encapsulation. The CPE runs IP on its LAN interface and PPP as the link layer protocol on the ATM WAN interface, as shown in Figure 2-18.

Figure 2-18 *PPPoA Network Cross Section*

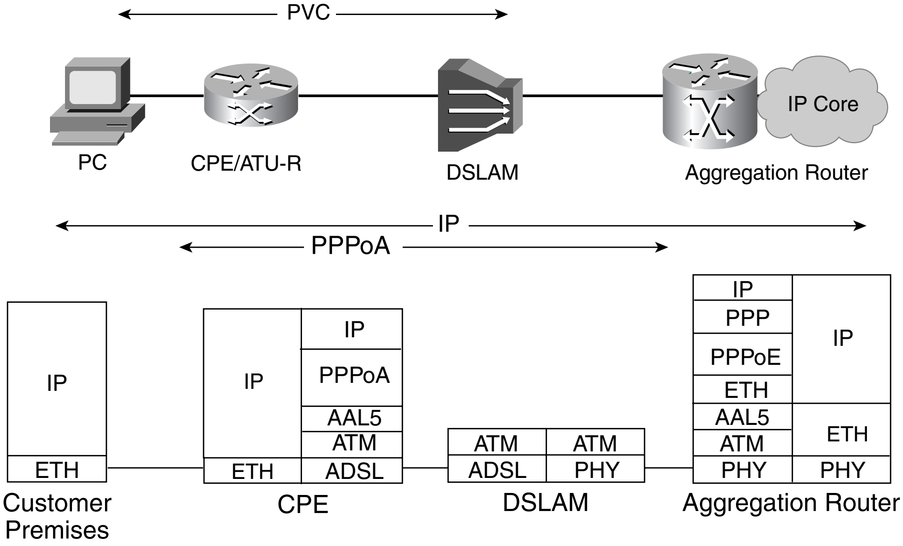

As soon as the PPP software detects that the ATM PVC is up, it tries to establish a session in classic PPP fashion. The router sends an LCP request, then changes state to authenticate and remains in this mode until either authentication succeeds or it times out. In case of timeout, the process starts over until a session is opened. This behavior can be disabled using **atm pppatm passive** but it is on by default on Cisco routers.

PPP Address Assignments

During PPP session negotiation, the client typically requests an address from the router. The router can find an address in one of several places. Either an **ip pool** is configured on the router, as in our example, or an address (or pool, for that matter) must be downloaded from a RADIUS server. (There is also the option of allowing the remote peer to keep its address, using **ip address negotiate** on the virtual template.)

The address download options are as follows:

- Download of a single address
- Download of a pool name
- Download of a pool at startup
- Use of On-Demand Address Pools

Download of a Single Address

When downloading a single address from RADIUS (using, for example, the **framed-ip-address** attribute), preference is given to downloaded attributes over parameters configured from the command line.

Example 2-16 shows how to use this attribute in a subscriber profile.

Example 2-16 Framed-IP-Address *RADIUS Profile*

```
jondoe Password = "cisco"
       Service-Type = Framed-User,
       Framed-Protocol = PPP,
       Framed-IP-Address = "192.168.11.1"
```

There are some special values for **framed-ip-address** that you should remember. If the AAA server returns a value of 255.255.255.254, that is an instruction to the router to fetch a dynamic IP address for this subscriber. In other words, the AAA server says to the router. "I want you to find a dynamic address for this subscriber." The router will probably do this using DHCP, and Example 2-17 shows the little bit of magic you need to combine DHCP and AAA.

Example 2-17 *Pools with DHCP*

```
! global command
ip dhcp-server 1.1.1.1
!
interface virtual-template 1
peer default ip address pool dhcp
```

Download of a Pool Name

In this scenario, a pool is already configured on the router and the RADIUS server just has to tell the router which one to use.

On the router, you would need the configuration in Example 2-18, with a **virtual-template** that *does not* have a pool name (or that would override the one coming from the RADIUS server) and a standard **ip address pool**.

Example 2-18 *Downloading Pool Names Router Configuration*

```
!Note no pool name
interface virtual-template 1
description PPPoE Clients
ip unnumbered loopback0
ppp authentication chap pap
ip local pool FOO 192.168.10.10 192.168.10.100
```

The corresponding AAA profile *with* a pool name would look like Example 2-19.

Example 2-19 *Downloading Pool Names RADIUS Profile*

```
janedoe Password = "cisco"
       Service-Type = Framed-User,
       Framed-Protocol = PPP,
       av-pair = "ip:addr-pool=FOO"
```

Download of a Pool at Startup

This trick is very similar to using AAA to download IP routes, in which case a Cisco IOS router can be configured to read static routing entries from an AAA server when it starts up. (The configuration is a little esoteric.)

The router is configured with a name that it uses to issue a RADIUS Access-Request at startup. The Access-Accept reply from the server includes one or more pools of IP addresses. These can be referenced in subscriber AAA profiles just as if they had been configured on the router. If a subscriber references an unknown pool, the router tries to download the complete list again. Example 2-20 shows how to activate this behavior in Cisco IOS. The router now sends an Access-Request using the name **load-pools**.

Example 2-20 *Downloading Pools Router Configuration*

```
! add this to the router configuration
aaa configuration config-username  load-pools
```

Now the RADIUS server needs a user profile that uses the same name as the router does, which is the case of the RADIUS profile in Example 2-21.

Example 2-21 *Downloading Pools RADIUS Profile*

```
nas1-pools Password = "cisco"
       Service-Type = Outbound-User,
       av-pair = "ip:pool-def#1=BAR 1921.168.11.10 192.168.11.100"
```

You can use these downloaded pools just like you do an IP address pool. Example 2-22 shows an example with a virtual-template configuration with a pool called BAR that was defined in Example 2-21.

Example 2-22 *Using Downloaded Pool Definitions*

```
interface virtual-template 2
description PPPoE Clients
ip unnumbered loopback0
peer default ip address pool BAR
ppp authentication chap pap
```

Use of On-Demand Address Pools

On-Demand Address Pools (ODAP), which is discussed in more detail in the Chapter 6, "Wholesale MPLS-VPN Related Service Features," is a powerful mechanism to allow a router to request IP pools dynamically from a DHCP server as existing ones are used.

All the techniques just described are well and good, but they let you allocate only a single host address (or pool from which a single address will be handed out to the host). What about the case of home networks with a broadband router connected to many different home computers? For that, the service provider needs other techniques, such as the following:

- **CPE PAT**—The CPE uses port address translation (PAT) to map many private host addresses to the single, public address assigned during PPP session negotiation.
- **IPCP subnet mask**—The CPE is assigned an IP subnet mask and an address during PPP session negotiation. It takes the first address of this subnet for its WAN interface and uses the rest of the pool to do network address translation (NAT) of host addresses. It is not possible to assign these addresses using DHCP, or else two links would be on the same IP subnet.

As usual, the subnet mask can either be configured on a b or in AAA. To make the IPCP subnet option work, you need to coordinate the aggregator, the CPE, and the AAA server.

Example 2-23 shows the configuration needed on the aggregation router.

Example 2-23 *IPCP Subnet Mask Router Configuration*

```
interface Virtual-Template2
ip unnumbered Loopback0
no peer default ip address
ppp authentication pap chap
ppp ipcp mask 255.255.255.240
!
```

Example 2-24 shows the configuration on the CPE, which needs to ask for the subnet mask when it brings up its PPP session.

Example 2-24 *IPCP Subnet Mask CPE Configuration*

```
!
interface Dialer 0
ppp ipcp mask request
```

Finally, Example 2-25 shows the RADIUS profile, which has a regular **framed-ip-address**, but also a subnet mask in the **framed-ip-netmask** attribute.

Example 2-25 *IPCP Subnet Mask RADIUS Profile*

```
CPE Password = "cisco"
Service-Type = Framed,
Framed-Protocol = PPP,
Framed-IP-Address=192.168.2.1
Framed-IP-netmask=255.255.255.248
```

PPP Quality of Service

There are two things to understand regarding QoS on PPP links: what type of QoS is supported and how to provision it.

One detail needs clarification at the outset: IP QoS on PPP interfaces is not as complete as on other Layer 2 interfaces. Depending on the actual router you are using, you can have classification, marking, and policing but probably not queuing. Of course, PPP runs on top of another Layer 2 interface and you can use the complete range of ATM CoS for DSL subscribers, for example. In each case (i.e., Layer 2 under PPP, or Layer 3 on the interface itself), use the classic Cisco IOS commands that were covered earlier in the chapter.

Provisioning QoS for PPP is a little different. Think first about what layer you are going to use to do the QoS and what type of QoS you need. Is it the policing at the ATM layer? Or classification at the IP layer? IP QoS commands, such as the **service-policy** command in Example 2-2, should go under the **virtual-template** on the router. ATM QoS parameters go under the PVC block. As with most things in PPP, you can provision QoS at either of these layers in a RADIUS profile.

Dynamic Bandwidth Selection (DBS) is the name of the Cisco IOS feature that lets you set a subscriber's ATM CoS profile using RADIUS. The idea behind the name is that a service provider would define different policies for subscribers: one for basic Internet access, one for VoIP, and so on. If there are different profiles predefined in RADIUS, then all the subscriber needs to do is to connect with a new username to get the new and improved QoS on his circuit. Using DBS means that the service provider operations team doesn't have to configure anything in the network as customers change back and forth between different levels of QoS.

Example 2-26 shows how to enable the DBS feature in Cisco IOS under an individual PVC, followed by Example 2-27, which gives the specific DBS RADIUS attributes.

Example 2-26 *DBS Router Configuration*

```
interface atm0/0/0.5 point-to-point
 ip address 192.168.2.1 255.255.255.0
 pvc 1/100
  dbs enable
  encapsulation aal5snap
  protocol ppp virtual-template1
```

Example 2-27 *DBS av-pairs RADIUS Profile*

```
Cisco-Avpair = "atm:peak-cell-rate=155000"
Cisco-Avpair = "atm:sustainable-cell-rate=155000"
```

The AAA parameters in Example 2-26 set the peak cell rate, which is mandatory and a sustainable cell rate. It behaves as UBR if only the PCR is given; otherwise it operates as a VRB-nrt circuit.

There is a similar set of special RADIUS attributes that let you download IP policing parameters to the aggregation router.

All in all, PPP is a little blunt when it comes to QoS support. With bridged access, there are very clear ways to map QoS policy between the IP layer and the transport layers, all of which have a good level of native QoS. The extra PPP layer blinds the access network, which cannot look into the PPP packets to know what QoS level to apply to the frames: There is no QoS marking in the PPP header and the original IP header is too deeply encapsulated to be able to look for the DSCP settings in hardware. In fact, all you are really doing with PPP QoS is prioritizing traffic on the aggregation router itself. None of the devices downstream (i.e., those between the aggregator and the subscriber) can automatically change its CoS settings if a subscriber uses a different DBS profile: This is not end-to-end QoS. Bridged access does offer, or is closer to offering, end-to-end QoS.

DSL Forum has been especially active in working on a new model that supports true multiservice traffic throughout the network, not just on the aggregation router. Interested readers should look for WT-59–related documents on the DSL Forum web site at http://www.dslforum.org.

PPP Authentication, Accounting, and Security

PPP has excellent authentication and accounting support. Millions of broadband and dial-up customers around the world use PPP for their Internet connection. The beautiful thing with PPP is that subscriber configuration can be centralized on a RADIUS server, which is a much more scalable way to run a network than to have to configure the devices independently. PPP is very well documented in other books, so the details are not going to be covered here, with the following two exceptions:

- PPP port-based authentication
- PPP security

Port-Based Authentication

Configuring a username in AAA for the CPE might seem easy to do but is in fact very awkward, because it means having a different username configured for each and every CPE. If a subscriber changes CPE, the username would have to be updated on the new device to make sure that the subscriber still gets the correct IP address. Rather than go to the pain of maintaining such a database, wouldn't life be easier if a CPE could be authenticated using the subscriber circuit ID? Happily, this is possible. The syntax on the router is **radius-server attribute nas-port format d**.

This simple statement makes the router include the circuit ID in the Radius NAS-port field. The format for ATM, which makes this value globally unique, is IP address/module/port/VPI/VCI. A corresponding format exists for Ethernet and VLANs (even QinQ) also. Obviously, the RADIUS server must support this.

PPP Security

Bridged access security is complex because it involves many subscribers who are all part of the same broadcast domain. Regardless of the actual tricks, DSL, cable, and Ethernet networks have many different bells and whistles to limit the broadcast domain as much as possible, ideally to a single subscriber.

PPP architectures just don't have this problem, because the subscriber links are actually routed interfaces and the aggregation router knows which address it assigned to whom. This removes a lot of the risk of IP and MAC layer spoofing, especially of the variety that lets one subscriber attack their neighbor *because of weaknesses in the aggregator or the broadband architecture itself.* It's important to be realistic here: suitably motivated subscribers at the other end of a PPP session *can* launch DoS and other nasty attacks. However, because the architecture provides a point-to-point link for each and every subscriber, there is inherently more security than on a network in which subscribers share the same Layer 2 segment.

You should remember the basic best practices for securing PPP connections and use CHAP authentication for the actual session itself. Also, protect the aggregation router. A PPP subscriber can still mount an attack against the default gateway. Ironically, PPP isn't as well served as Ethernet (but it doesn't have the same risks of ARP- and broadcast-based attacks), but URPF and NetFlow are also really good techniques to use in PPP architectures. A PPP-specific attack would be to launch a DoS attack against the RADIUS server by opening a gazillion sessions, or opening and closing them constantly. Cisco IOS can limit the number of sessions per connection or per MAC address, and this is a good feature to turn on. Perhaps best of all, RADIUS billing records provide a great way to track unusual usage back to an individual subscriber—even if he spoofs IP addresses and blasts a week's worth of traffic in a couple of hours, his billing record will show the volume of traffic sent over the subscriber line (ATM VC or Ethernet port).

Summary

This chapter reviewed the access technologies and protocols for DSL, cable, and Ethernet broadband networks. Even though there are major differences in the physical network, at Layers 2 and 3, they are much more similar. There are, in effect, only so many ways to connect a broadband customer, and a solution must offer efficient mechanisms for authentication, accounting, routing, and so on. Bridged access is prevalent on cable and Ethernet. PPP is more common on DSL but is used on the other two access types, because they both have an Ethernet MAC layer.

Bridged access is very simple and cost effective. It takes advantage of the fact that everyone is using Ethernet these days, and the bridged CPE is as cheap as they come. DHCP is a tried and trusted mechanism for managing addresses dynamically, and there are some enhancements in Cisco IOS that let ISPs hand out individual addresses to subscribers as they connect, but announce aggregate routes over their core networks.

PPP is very complete and lets service providers centralize customer configuration in one place. When a subscriber connects using PPP, the aggregation router authenticates using RADIUS and retrieves the subscriber's configuration from the central server. Operationally, this is a huge advantage. PPP offers many different ways to manage addresses, but they can be summarized as either allocating single addresses to individual subscribers, or allocating blocks of addresses to aggregation routers.

Security is always a concern on bridged networks and there are different solutions found on the different types of broadband-access platforms. DSL really cheats a little, and RBE treats bridged subscribers as if they were on routed, point-to-point links. The advantage to doing this is that broadcast domains are not shared across multiple subscribers. Cable has a sophisticated Layer 1, called DOCSIS, as well as enhancements on the cable modem and CMTS that address the same problems. Finally, Ethernet has a lot more security mechanisms than people usually expect, and switches offer protection that helps against spoofing and DoS attacks. PPP security is less of a major issue because there is no concept of a shared broadcast domain that works against you.

Address management and routing are really similar across all the architectures. PPP and bridging use different control protocols to allocate addresses, but once they have done so, the same guideline applies: Summarize routes early.

End-to-end QoS is much easier to achieve using a bridged architecture. Layer 2 to Layer 3 QoS mapping is well understood, and some of the broadband networks have rich native QoS capabilities—DSL because of ATM, and cable with DOCSIS. By and large, aggregation routers also allow service providers to combine this with the full gamut of Layer 3 QoS. The net result is that broadband services can carry multiservice traffic, such as Voice over IP, correctly. QoS support is a weakness of the PPP architecture. Different, even dynamic, QoS profiles are supported on the aggregation router, but end-to-end QoS across the access network is very hard to achieve.

Chapter 3 moves from the access part of the network to the service end and looks at the three major types of VPN technology used with broadband access today: GRE, IPSec, and L2TP.