

RFC1483 Bridging Baseline Architecture

Document ID: 12916

Contents

Introduction

Assumption

Technology Brief

Advantages and Disadvantages of RFC1483 Bridging

Advantages

Disadvantages

Implementation Considerations

Network Architecture

Design Considerations

Key Points of this Architecture

How a Service Destination is Reached

Operational Description

Conclusion

Related Information

Introduction

This document describes end-to-end asymmetric digital subscriber line (ADSL) architecture when using RFC1483 bridging. Note that most early versions of xDSL modems were bridges between 10BaseT Ethernet at the host side and RFC1483 encapsulated bridge frames on the WAN side. Even today, the majority of the ADSL customer premises equipment (CPE) deployed in the field are in pure bridging mode.

Assumption

The baseline architecture is designed with the assumption of providing high speed Internet access to the end subscriber using the RFC1483 bridging model and ATM as the core backbone. The content of this document is based on the architecture of existing deployments and some inhouse tests.

Technology Brief

RFC1483 describes two different methods for carrying connectionless network interconnect traffic over an ATM network: routed protocol data units (PDUs) and bridged PDUs.

Routing allows multiplexing of multiple protocols over a single ATM virtual circuit (VC). The protocol of a carried PDU is identified by prefixing the PDU with an IEEE 802.2 Logical Link Control (LLC) header.

Bridging performs higher-layer protocol multiplexing implicitly by ATM virtual circuits. For more information, refer to RFC1483.

This document refers only to bridged PDUs.

Advantages and Disadvantages of RFC1483 Bridging

Following is a summary of the advantages and disadvantages of the RFC1483 bridging architecture. This architecture has some important disadvantages, most of which are inherent in the bridging model. Some of the

disadvantages were noticed during ADSL deployments at customer sites.

Advantages

- Simple to understand.

Bridging is very simple to understand and implement because there are no complex issues such as routing or authentication requirements for users.

- Minimal configuration of the CPE.

The service provider considers this important because it no longer requires a large number of truck rolls and no longer needs to invest heavily in personnel for the support of higher level protocols. The CPE in bridge mode acts as a very simple device. Minimal troubleshooting is involved at the CPE because everything that comes in from the Ethernet passes directly to the WAN side.

- Easy to install.

Bridging architecture is easy to install because of its simplistic nature. After end-to-end permanent virtual circuits (PVCs) are established, activities such as IP at the upper layer protocols become transparent.

- Multiprotocol support for the subscriber.

When the CPE is in bridging mode, it is not concerned with which upper layer protocol is being encapsulated.

- Ideal for Internet access in a single user environment.

Because the CPE acts as a set-top box, complex troubleshooting is not required for upper layer protocols. The end PCs do not require additional client installation.

Disadvantages

- Bridging depends heavily on broadcasts to establish connectivity.

Broadcasts between thousands of users are inherently unscalable. The reasons for this are that the broadcast consumes bandwidth across the users' xDSL loop, and the broadcast requires resources at the head-end router to replicate packets for the broadcast over point-to-point (ATM PVC) media.

- Bridging is inherently insecure and requires a trusted environment.

The Address Resolution Protocol (ARP) replies can be spoofed and a network address hijacked. Additionally, broadcast attacks can be initiated on the local subnet, thus denying service to all members of the local subnet.

- IP address hijacking is possible.

Implementation Considerations

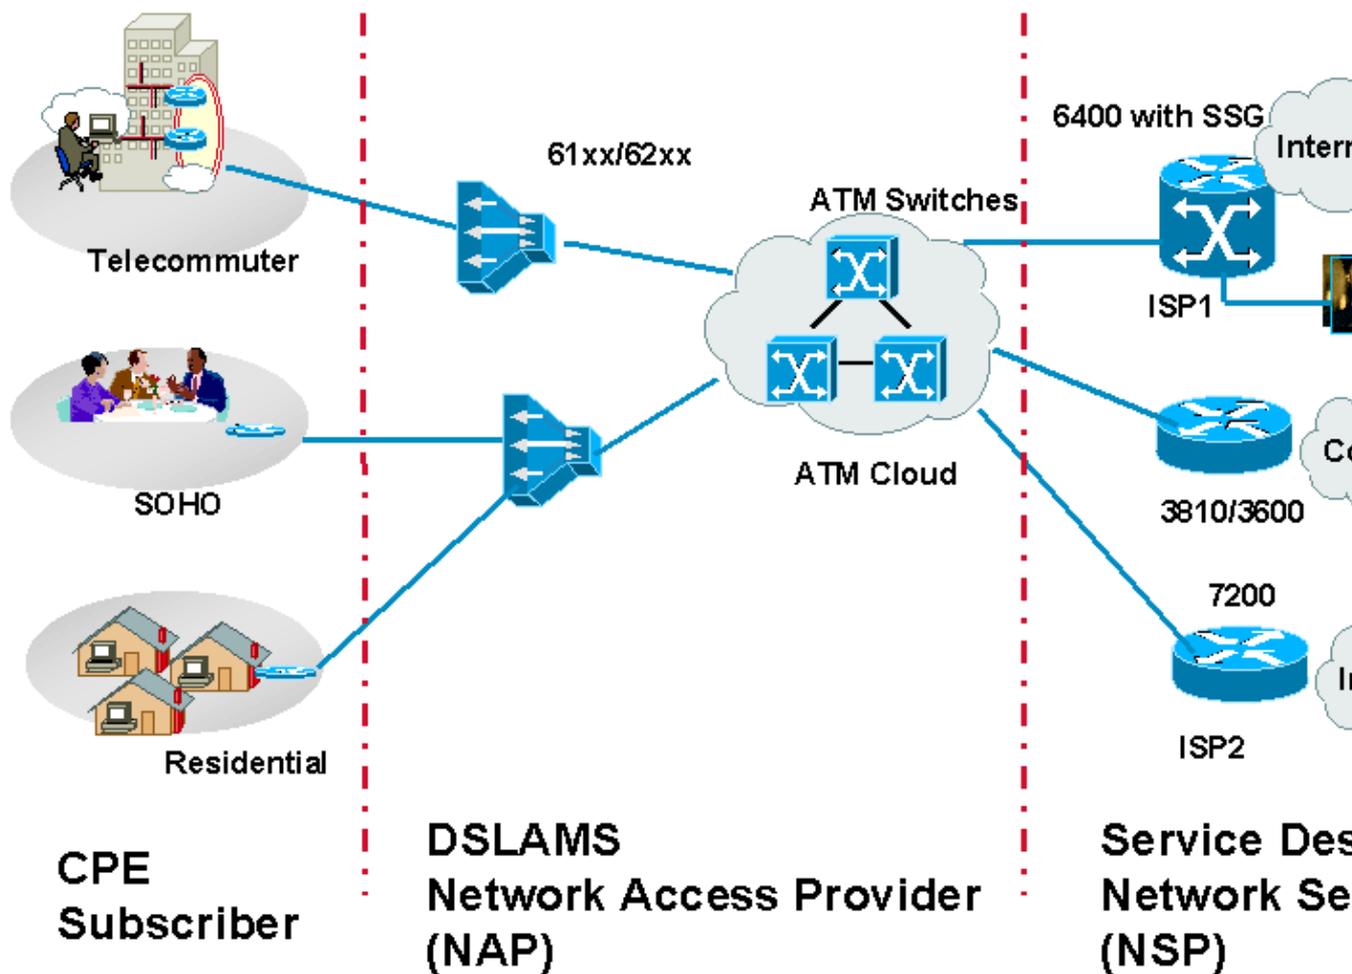
Consider the following questions before implementing the RFC1483 bridging architecture.

- What are the current and planned numbers of subscribers to be serviced?
- Do the subscribers need to communicate with each other?
- Are these subscribers single-user residential customers? Do you service small office, home office (SOHO) customers who might have a small LAN behind the CPE?
- What is the deployment and provisioning of CPEs, digital subscriber line access multiplexers (DSLAMs) and aggregation Post Office Protocols (POPs)?

- Are the Network Access Provider (NAP) and the Network Service Provider (NSP) the same entity? Does the business model for the NAP also involve selling wholesale services such as secured corporate access, and value-added services such as voice and video?
- Does the NSP want to offer service selection capabilities?
- How can the accounting and billing be achieved? Is it per usage, per bandwidth, or per service?
- Is the business model of the company that of an independent local exchange carrier (ILEC), competitive local exchange carrier (CLEC), or Internet Service Provider (ISP)?
- What types of applications does the NSP want to offer to the end subscriber?
- What is data flow volume both upstream and downstream?

With these points considered, following are descriptions of how the RFC1483 bridging architecture will fit and scale to different business models.

Network Architecture



RFC1483 Bridging: Network Architecture

Design Considerations

As previously mentioned, there are some inherent problems with the RFC1483 bridging architecture.

The IOS subscriber bridging feature addresses some of these problems. Selective application of subscriber policies to a bridge group controls the flooding of ARPs, unknown packets, and others down each ADSL loop. For example, by preventing ARPs from being broadcasted, a hostile user cannot discover the IP address of another user.

Another solution is to put all subscribers into a single subinterface. Normal bridging behavior will not forward frames to the port on which the frame was received. In essence, this enforces a type of subscriber bridging in

which all packets between subscribers are filtered. However, this approach has the following flaws:

- Subscriber policy is only applied between subinterfaces. To apply subscriber policies between two different users, each user must be in a different ATM subinterface.
- Since the Layer 2-to-Layer 3 address mapping is learned (via ARP), hostile users can still hijack the connection of other users. This is done by generating ARP traffic with another user's IP address and using a different MAC address.

The second scenario is more serious for the carrier or ISP. In this situation, any user can assign the wrong address to a PC or Ethernet-attached device such as a printer, and cause connection problems for another user. Such errors or attacks are hard to pinpoint and correct because the offender can be tracked only by tracing the MAC address of the offender.

Some carriers try to work around this problem by segregating users across bridge groups, and by implementing subscriber bridging across subinterfaces. In this case when integrated routing and bridging (IRB) is required, each user is assigned a unique bridge group and Bridge Group Virtual Interface (BVI). This approach uses two interfaces per subscriber and can be challenging to manage.

These issues are addressed and resolved in some ways by the Routed Bridged Encapsulation (RBE) feature which was introduced in Cisco IOS® Software Release 12.0(5)DC on the Cisco 6400.

Considering some of the disadvantages of bridging, you might wonder why the bridging architecture would ever be implemented. The answer is simple. Most of the ADSL CPEs installed in the field are capable only of forwarding bridged frames. In these cases, the NSP must implement bridging.

Today, CPEs can do Point-to-Point Protocol over ATM (PPPoA), RFC1483 bridging, and RFC1483 routing. The NSP determines whether to do bridging or PPP. The decision is based on the implementation considerations mentioned earlier, in addition to the pros and cons of each architecture.

Even with the disadvantages of bridging architecture, it may be suitable for a small ISP (that may not be the NAP) or an NAP/NSP serving a smaller number of subscribers. In these scenarios, the NAP usually forwards all subscriber traffic to the ISP/NSP, which terminates those subscribers. The NAP could choose to provide subscriber traffic using ATM or Frame Relay as the Layer 2 protocol.

The NAPs using current generation DSLAMs can only transport subscriber traffic using ATM. In this case, the ISP should terminate ATM permanent virtual circuits (PVCs) to a router.

If the ISP/NSP does not have the ATM interface, a regular serial interface with encapsulation ATM Data Exchange Interface (DXI) (possibly on an additional device) can be used to accept the incoming bridged PDUs.

In both scenarios, the NSP/ISP may have to configure IRB on the router (except when using encapsulation ATM DXI or in the case of transparent bridging). Today, the most common practice for terminating bridged subscribers on the NSP/ISP router is to implement IRB. (It is expected that service providers will gradually migrate to RBE.)

Because of some of the limitations mentioned above, the NSP/ISP may opt to configure separate bridge groups for each set of subscribers or to configure all the subscribers in one bridge group. The common practice is to configure a few bridge groups, and then configure all the subscribers under separate multipoint interfaces. As mentioned earlier, the subscribers under the same multipoint interface may not be able to communicate with each other. If certain users need to communicate, configure those subscribers under different interfaces (they can still be in the same bridge group).

For a small ISP/NSP, the most common routers used to terminate bridged subscribers are the Cisco 3810, Cisco 3600, and Cisco 7200. For an ISP/NSP with a large subscriber base, the Cisco 6400 is preferred. Before calculating the memory requirements for these routers, consider the same factors as for any other environment: number of users, bandwidth, and router resources.

Key Points of this Architecture

Following are the key points of the architecture.

CPE

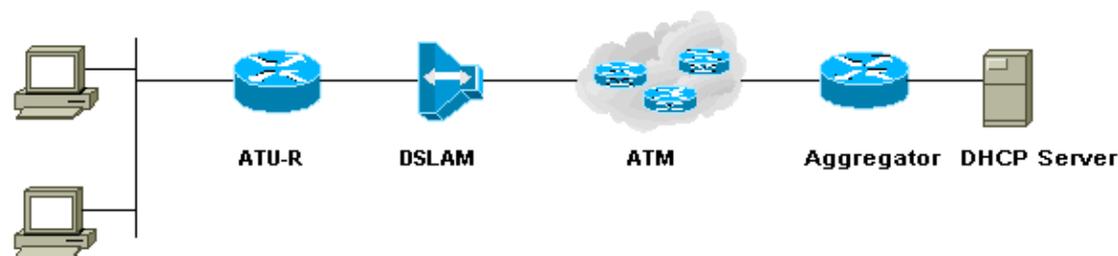
Cisco offers various CPEs that operate with Cisco and non-Cisco DSLAMs. The configuration for each of these CPEs is problem free and requires no input from the subscriber. The primary requirement is that the CPE define an ATM virtual path identifier/virtual channel identifier (VPI/VCI). This allows the CPE to train up with the DSLAM and start passing traffic. In most instances, the NAP opts to configure the same VPI/VCI for all subscribers. The NAP usually pre-provisions the CPE before deploying it at the subscriber's location.

In bridging architecture, the main consideration for the CPE and its deployment is how the NAP will manage the CPE after it is installed in the field. This is a concern because bridging does not require an IP address for the CPE. However, Cisco CPEs can be provisioned with an IP address in bridging mode. The NAP may use this feature to Telnet to the CPE to collect statistics or to help the subscriber with troubleshooting. To allow CPEs to be managed through the DSLAMs, new proxy element functionality is being added.

In bridging mode, if no management IP address is assigned to the CPE, the operator can only manage the CPE through the CPE management port. If a management IP address is assigned, the operator can use a Hypertext Transfer Protocol (HTTP) browser to manage the device. However, this option is generally not available.

When the CPE is in bridging mode, the service destination (which could be the NSP/ISP) should provide an IP address that will be used as the default gateway for the PCs behind the CPE. These PCs must be set to the correct default gateway. Otherwise, even if the modem is trained (which means that the physical layer is good between the CPE and the DSLAM), the subscriber may not be able to pass traffic. This is not an issue if Dynamic Host Configuration Protocol (DHCP) is used to assign subscriber DHCP addresses because the default router is returned by the DHCP server.

IP Management



RFC1483 Bridging: IP Management

In a bridged environment, the IP addresses are allocated to the end stations by a DHCP server located at the service destination, usually in the NSP/ISP network. This is the most common approach and is implemented by most NSPs/ISPs using this model.

Another approach is to provide static IP addresses to the subscribers. In this case, either a subnet of IP addresses or a single IP address is allocated per subscriber, depending on the subscriber requirements. For example, subscribers wanting to host a Web server or an email server will need a set of IP addresses rather

than a single IP address. The problem with this is that the NSP/ISP has to provide public IP addresses and may quickly run out of them.

Some NSPs/ISPs have provided private IP addresses to their subscribers. They then perform Network Address Translation (NAT) at the service destination router.

NSPs/ISPs that provide a full subnet for one bridge group (with more than one subscriber) should know that one user can assign the wrong address to a PC or Ethernet–attached device, such as a printer, and cause connection problems for another user.

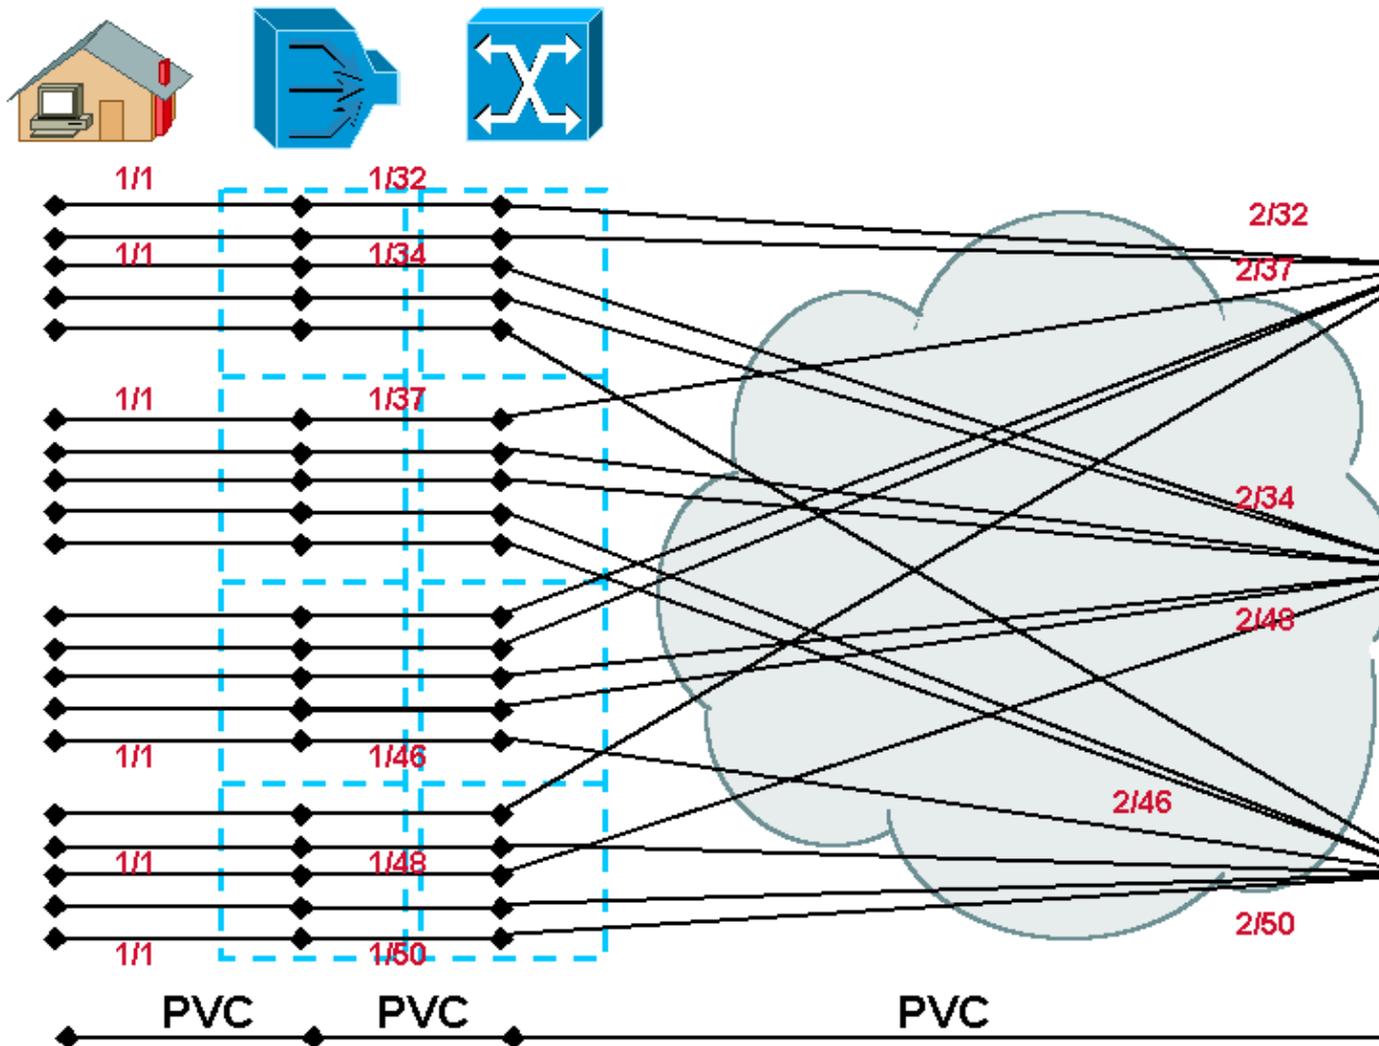
It is also possible for an NSP/ISP to restrict the number of PCs that can access the service at one time. This is done by configuring the maximum users on the Ethernet interface.

However, this method has the following flaw. If three PCs are configured to use the service and one of the subscribers adds a network printer (which has its own MAC address) during a time when one of the PCs is idle, the PC's MAC address will disappear from the ARP entry of the CPE.

If the printer becomes active while a PC is idle, the printer's MAC address will be entered in the ARP entry. When a user decides to use this PC to access the Internet, it will be unavailable because the CPE already has allowed three MAC entries. The strategy of limiting users on the CPE can be used, but care should be taken in fixing the numbers.

How a Service Destination is Reached

End-to-End PVC



RFC1483 Bridging: End-to-End PVC

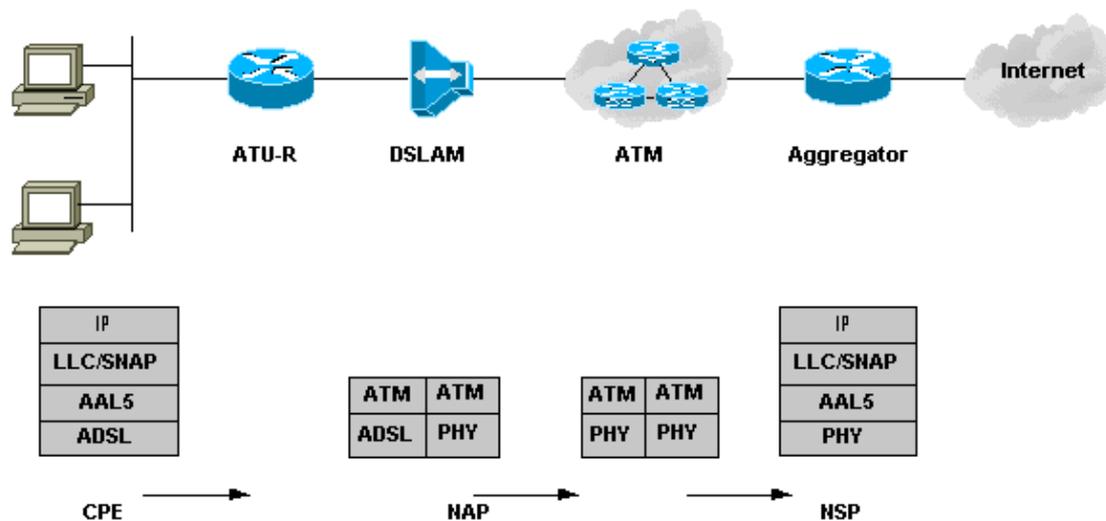
In an end-to-end PVC architecture with bridging, the service destination is reached by the creation of PVCs between each hop. However, management of these PVCs can be challenging for the NAP/NSP. Additionally, the number of PVCs that can be defined through the ATM cloud is limited. This limitation affects many of the NAPs/NSPs who adopt an end-to-end PVC model. For each subscriber there will be a fixed, unique set of VPIs/VCIs along the entire path. Switched virtual circuits (SVCs) help to overcome some of these problems, and many access providers are migrating to IP-enabled core networks to solve the problem of VC exhaustion.

The NSP/ISP also has the option of using the Cisco Service Selection Gateway (SSG) functionality to provide different services to subscribers.

In this architecture, secured access to a corporate gateway is achieved by terminating the subscriber traffic PVC straight in the corporate router at Layer 2. The PVC-based architectures are inherently secure when

sharing data with other service destinations.

Operational Description



RFC1483 Bridging: Operational Description

The Cisco 6xx CPE defaults to routing mode. Therefore, when it is configured for bridging mode and installed at the subscriber's location with the necessary splitters/microfilters, it trains up automatically upon power up. When the CPE trains up, it indicates that the physical layer between the CPE and DSLAM is fine. Depending on how the end station's IP address is configured (that is, whether it is assigned via a DHCP server or it is a static IP address with default gateway information), it can then communicate with the service destination.

Following is a description of the flow of packets.

The user's data is encapsulated in IEEE 802.3 from the PC and enters the Cisco 6xx CPE. It is then encapsulated into an Logical Link Control/Subnetwork Access Protocol (LLC/SNAP) header, which in turn is encapsulated in ATM adaptation layer 5 (AAL5) and handed over to the ATM layer.

The ATM cells are then modulated by the ADSL transmission technology, Carrierless Amplitude and Phase (CAP) modulation or Discrete Multi-Tone (DMT), and sent over the wire to the DSLAM. At the DSLAM, these modulated signals are first received by the POTS splitter, which checks whether the frequency of the signal is below or above 4 kHz. After it identifies the signals as above 4 kHz, it passes them to the ADSL Transmission Unit – Central Office (ATU-C) in the DSLAM.

The ATU-C demodulates the signal and retrieves the ATM cells, which are then passed to the network interface card (NIC) in the multiplexing device (MUX). The NIC looks at the subscriber side VPI/VCI information in the ATM header and makes the switching decision to another VPI/VCI which will be forwarded to the service destination router. After the service destination router receives these cells on a particular ATM interface, it re-assembles them, looks at the upper layer, and passes the information to the BVI interface. The BVI interface looks at the Layer 3 information and decides where the packet is to be delivered.

Conclusion

The RFC1483 bridging model is more suitable for smaller ISPs or corporate access for which scalability does not become an issue. Because it is very simple to understand and implement, it has become the choice of many smaller ISPs. However, as a result of some security and scalability issues, bridging architecture is losing

its popularity. NSPs/ISPs are opting for RBE or moving toward PPPoA or PPPoE, which are highly scalable and very secure, but more complex and difficult to implement.

Related Information

- **DSL Technical Support**
 - **Technical Support – Cisco Systems**
-

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2014 – 2015 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: Nov 23, 2007

Document ID: 12916
