An Overview of Point to Point Protocol Architectures in Broadband Access Networks

Pranesh Kulkarni
Ixia Technologies Pvt. Ltd.
Bangalore, India
pkulkarni@ixiacom.com
www.ixiacom.com

Shantala Patil
Dept. of Computer Science and Engineering
Reva Institute of Technology and Management
Bangalore, India
patilss82@gmail.com
www.revainstitution.org

***Abstract*—Access technology holds key to getting attached to Internet from customer premises equipment in a broadband network. The use of different physical media types such as ATM, Ethernet and SONET have further paved way for use of one or combination of such access protocols to each one of media types. However, fundamentally Point to Point Protocol (PPP) has emerged as predominant solution for accessing broadband networks, and has generated wider acceptability among service providers. This paper discusses broadband network architecture and how PPP fits in overall scheme of things in broadband access network. Further various PPP protocol extensions over Ethernet (PPPoE), ATM (PPPoA) and combination (PPPoEoA) have been individually discussed in the context of their use cases.**

I. INTRODUCTION

For millions of users across the world, Internet access has been viewed as not mere a luxury, but an absolute necessity. Given the volume of the information, especially the multimedia data that is accessed by these users, there gets created scope for wider access bandwidths. This led to advent of faster data rate broadband access, which is discussed briefly in section II. In section III we discuss PPP, which Service providers started using in the early 1990s to connect telephone dial-up and ISDN subscribers to the Internet. Designed to connect a wide variety of hosts, bridges, and routers, PPP can multiplex different higher-layer application protocols simultaneously over the same link. Today, service providers continue to use PPP, connecting multiplay subscribers through high-speed broadband DSL. In the last section we discuss three prominent network architectures based on two fundamental media types, Ethernet and ATM. These are being widely deployed in the broadband access networks by service providers.

II. BROADBAND ACCESS NETWORK

A. *Overview*

Figure 1 depicts the basic broadband network architecture.

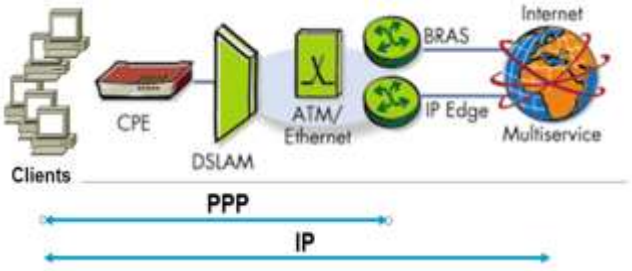

Figure 1. Broadband Access Network

The twisted pairs running from customer premises from multiple locations are grouped together with a device known as a DSL Access Multiplexer (DSLAM) which combines all of the individual user sessions onto an ATM/Ethernet trunk. Next, the traffic from multiple DSLAMs is further combined at an aggregation point, typically located in a central office. From there the traffic is forwarded to a router, which in turn connects to the Internet. In a dial-up environment, PPP sessions are terminated on a Remote Access Server (RAS). Broadband Networks use a similar device known as a Broadband Remote Access Server (BRAS) to terminate large quantities of PPPoA connections. This device is also commonly called an aggregation router or aggregator and is located in a central office facility of a service provider.

III. POINT TO POINT PROTOCOL

Point-to-point Protocol (PPP) [1] is used for communications between two nodes, such as between a client and a server. Service providers started using PPP in the early 1990s to connect telephone dial-up and ISDN subscribers to the Internet. Designed to connect a wide variety of hosts, bridges, and routers, PPP can multiplex different higher-layer application protocols simultaneously over the same link. Today, service providers continue to use PPP, connecting multiplay subscribers through high-speed broadband DSL.

A. *PPP in Context of Broadband Access Networks*

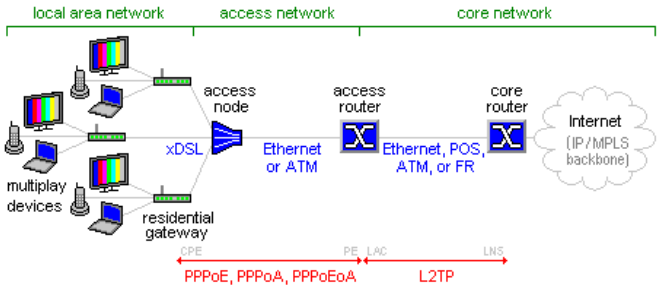

Figure 2. PPP in the Context of Broadband Access Networks

Figure 2 shows one way in which residential devices access the Internet. In this scenario, the:

- multiplay devices—include computers, IPTV set-top boxes, and IP phones used by a subscriber. These devices connect via Ethernet to a residential gateway.
- residential gateway—is a DSL router that:
    - serves as a PPP client or CPE, establishing a PPP link with an access router
    - configures IP addresses for multiplay devices in its LAN, and forwards traffic to and from the multiplay devices
- xDSL—is a DSL service such as ADSL or HDSL running over POTS
- access node—aggregates the data from multiple DSL lines to a single access router, through a high-speed Ethernet or OC-3/12 ATM network
- access router—
    - terminates multiple PPP sessions and forwards subscriber traffic to a core router, through a higher-speed GbE or OC-12 POS/ATM/Frame Relay link
    - serves as the PE end of a PPP link
    - serves as a LAC, using L2TP to extend PPP sessions by tunneling them to an LNS
- core router— forwards subscriber traffic through the Internet core using IP routing tables, MPLS, VRF routing, ATM PVCs, etc.
    - serves as an LNS, terminating PPP sessions initiated by a LAC and forwarding subscriber traffic through the Internet core to the remote destinations

A PPPoX session provides IP addressing, user authentication, billing, and QoS for each PPP client. A session can run over a variety of link layers, for example:

- PPPoE—PPP over Ethernet [2]
- PPPoA—PPP over ATM [3]
- PPPoEoA—PPP over Ethernet over ATM
- L2TP—PPP over L2TP

PPP which was originally defined for a direct connection between devices over a leased line using, several methods have been defined to establish PPP connections across other media. These include PPP over ATM (PPPoA), PPP over Ethernet (PPPoE), and PPP over SONET/SDH (POS).

Hence, there are two different forms of this protocol commonly implemented in broadband access networks: PPP over Ethernet (PPPoE) and PPP over ATM (PPPoA). Additionally, there also is a combination of both of these protocols called PPP over Ethernet over ATM or PPPoEoA. PPP offers all of the flexibility and ease of use associated with the other architectures and adds a level of security as well. Hence, it is the protocol of choice for most broadband service providers.

In the following sections we shall primarily discuss PPP architecture over Ethernet (PPPoE) and ATM (PPPoA).

B. *History of PPP development*

PPP is ancient according to Internet technology – it was initially adopted as a standard in 1989 (RFCs 1171 and 1172). It has been updated several times since then and its most recent version is contained in RFC 1661 from July 1994. PPP was designed to support communications between devices (typically routers) over traditional telco leased-line circuits (i.e., T1s). Since line quality was a major concern in the late 1980s and early 1990s, there are several mechanisms intrinsic to PPP to help it tolerate circuit degradation. In order to facilitate this adaptation, PPP was designed to be a connection-oriented protocol. This would appear to be unnecessary since it was initially built for use on point-to-point circuits which can only have two unique end-points. In addition to link quality checks (the Link Control Protocol (LCP) is a subset of PPP which is used for this purpose), the fact that PPP is connection-oriented also enhances security by requiring that packets are forwarded only to specific, known addresses.

In the early 1990s, PPP was adapted to support dial-up links. Again, being connection-oriented, it helped secure the links. Furthermore, since a “handshake” procedure is required as the link is initiated; it was easy to insert an authentication algorithm into this process. In a dial-up environment, a device known as a Remote Access Server (RAS) will validate the users ID and password. PPP will then provide for the automatic configuration of a link, including supplying the remote device

with an appropriate IP address. All of these capabilities have combined to make PPP the favorite protocol for dial-up connectivity – PPP is supported by virtually all ISPs. In fact, dial-up PPP is included in most commonly used PC operating systems.

C. PPP Session Establishment

PPPoE includes a straightforward mechanism for the host to find PPPoE server/BRAS to communicate with. The host broadcasts a request to establish a connection (PADI); all potential BRAS devices respond (PADO) with an "offer" to be the termination point; the host selects one (PADR); and the BRAS responds by assigning a session identifier (session-id). PPPoE flows also typically include the PPP link establishment phase. Originally used to establish the dial-up connection, this phase negotiates line characteristics such as the maximum MTU size, the authentication protocol to be used, and the link quality monitoring protocol to be used.

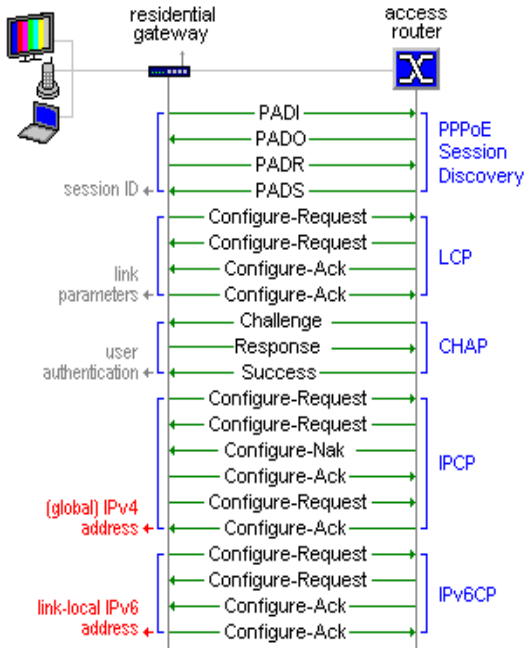

Figure 3. PPP link setup

1) Subscriber Authentication:

PPP authenticates users before allowing them access to the network, typically by requiring that the user log into the network using an assigned userid and password.

PPP authentication is tightly integrated with RADIUS. During this authentication phase, the network assigns attributes to individual subscribers by forwarding the login request to a RADIUS server. The RADIUS server returns information that allows the BRAS to determine what to do with the session (filters, multicast enable/disable, bandwidth control, QoS control, policy routing rules, L2TP LNS destination, etc.). Once the session gets successfully established, RADIUS accounting will start which allows to provider to do either time-based or volumebased billing. When the session gets disconnected, either explicitly or by means of missing keepalives, RADIUS accounting for the session will be closed by means of RADIUS accounting stop messages.

Offloading the session establishment, control, and accounting to RADIUS simplifies subscriber management. Subscriber-specific data is kept offline in the RADIUS system and the subscriber sessions are created dynamically during PPP LCP and IPCP transactions [4]. The BRAS acts as an intelligent mediation device between the client, AAA, and policy system while acting as the IP control point at the network edge.

These aspects of PPP enable subscriber control in an automated and centralized fashion. The session concept with its short keepalive interval allows the service provider to have accurate accounting data for each subscriber and maintain an accurate view on when subscribers are actually online which could be important for legal purposes. The tight integration with RADIUS allows for centralized policy and QoS control plus detailed accounting information on a per subscriber basis, so whenever services change there is no need to touch any access equipment or to reprovision any subscriber's connection. Enabling lawful interception on a per subscriber basis can now also completely enabled via RADIUS making the RADIUS integration even more complete.

2) Address Assignment:

PPP incudes a process for assigning Layer 3 attributes using network control protocols (NCPs). The NCP used to assign IP addresses within a PPP connection is IP Control Protocol (IPCP) [4].

The separation between link establishment and IP address allocation makes it possible to figure out who the subscriber is before deciding how to treat his session: Either terminate the session locally or tunnel the session to a wholesaler. The domain name provided in the user's credentials enables automated service or ISP selection. This powerful aspect of PPP is a key reason dictating why PPP will remain the predominant session protocol for providers that are offering wholesale services to third party ISPs.

3) Session Monitoring:

Another important aspect of PPP is that it is a session-based protocol which monitors line quality. Using PPP keepalives, both endpoints can monitor whether the session is still up and running. Typical keepalive times are in the order of 30 seconds. Upon missing a few consecutive keepalives, the BRAS will terminate the session and clean up all state information. The client will typically try to re-establish the session automatically. In case of redundant BRAS setups (more typical for Ethernet-DSL than for ATM-DSL), the redundant BRAS could now start to accept incoming PPP sessions.

IV. PPP NETWORK ARCHITECTURES

A. PPPoA

PPP over ATM (PPPoA) [3] was the connection method originally specified by the DSL Forum, and is the most prevalent method for connecting broadband users into the network. At an intermediate point such as a DSLAM or edge router, the individual PPPoA sessions are aggregated onto a single ATM VC uplink. This allows the network to scale, since the backbone could not scale to support a unique ATM VC for each subscriber.

All of the advantages of PPP in a dial-up environment are also extensible for higher speed broadband connectivity. However, the commonly used broadband topologies (ADSL and cable modems) are not based upon point-to-point circuits. Instead, they utilize some degree of network and/or component sharing to gain economies of scale. Hence PPP needed to adapt yet again to accommodate the evolving network architecture. Since ADSL is based largely on an ATM infrastructure, PPP over ATM (PPPoA) is used for this application [3]. This topology also lets service providers take advantage of ATM's sophisticated quality of service facilities to provide service level agreements (SLAs) to their customers.

In a dial-up environment, PPP sessions are terminated on a Remote Access Server (RAS). Broadband networks use a similar device known as a Broadband Remote Access Server (BRAS) to terminate large quantities of PPPoA connections. This device is also commonly called an aggregation router or aggregator and is located in a central office facility of a service provider. The user end of the PPPoA connection can be terminated at the user's workstation, but this will typically require the installation of additional complex software and an expensive ATM network interface card (NIC) in the computer. This complexity and expense are usually conveniently hidden from the user by installing the protocol stack in the ATU-R modem. Hence the PPPoA protocol usually operates between the DSL modem and the aggregator (a.k.a. BRAS).

In brief, PPPoA originates at the ATU-R at the user's location and terminates in the BRAS at the service provider's facility. An ATM permanent virtual circuit (PVC) carries the protocol (and the user data) between these two devices. The BRAS then aggregates all of the PVCs into an IP flow, which then is forwarded to upstream routers and the Internet.

B. PPPoE

As networks transition to Ethernet, PPPoE [4] has emerged as an alternative to PPPoA. There are several advantages to PPPoE. PPPoE can be implemented in software on PCs, while PPPoA requires a special ATM line card (or a DSL modem which supports PPPoA). In addition, PPPoA requires a separate ATM VC for each service, while PPPoE allows multiple services over the same connection.

The PPPoE usually involves typical ADSL broadband solution. Cable modems are supported by an even simpler architecture. Since the PC is connected to a cable modem via Ethernet, there is no need to translate multiple protocols. Instead, Ethernet can be extended across the cable network all the way to the cable modem equivalent of a central office. Therefore, it is logical to extend PPP the entire length of the connection thus increasing the range of the benefits associated with PPP and also eliminating the requirement for any protocol translation. Hence, it is feasible to run PPP on the workstation and continue that session all the way to an aggregation router at the ISPs facility.

The only drawback to a PPPoE solution is that it requires the addition of PPPoE software on the workstation. However, this shortcoming can also be turned around and used as an advantage for service providers. Service providers can distribute CDs with the appropriate PPPoE protocol stack (commonly called a "shim") to their customers. Additionally, they can include other programs such as a GUI front-end for their service, advertising, and other management and diagnostic tools. Thus the PPPoE client software can be installed transparently, along with whatever accompanying utilities the service provider desires.

C. PPPoEoA

In the ADSL model, it should be noted that an ATM PVC extends from the ATU-R to the BRAS. However, the user's workstation is connected to the ATU-R via the ubiquitous Ethernet protocol. The ATU-R functions as a bridge between the two protocols. Since the PPP protocol stops at the ATU-R, the user's Ethernet segment of the connection is not subject to the benefits of the PPP Link Control Protocol. Therefore, when the user's PC is turned off, the service provider may still believe that a valid connection exists since the LCP keep-alive messages will continue to be answered from the powered-on ATU-R. This effects network security and also usurps network resources that could otherwise be used by active customers. This problem can be resolved within a DSL environment by extending PPP to the workstation in a manner comparable to the cable modem architecture. Initially this could only be accomplished by installing an expensive ATM NIC in the PC. However, a more progressive alternative is to invoke the PPP over Ethernet over ATM (PPPoEoA) protocol. PPPoEoA is not an approved standard within any of the usual standards bodies. However, it is a de facto standard that is growing in popularity throughout the broadband service provider community. Simply put, PPPoE is enabled on the workstation. This is then encapsulated in ATM at the ATU-R for transit over the service provider's network. The aggregation router then strips off the overhead bytes and is left with a valid IP datagram that can be forwarded to the Internet.

PPPoEoA which involves multiprotocol encapsulation over ATM Adaptation Layer 5 [5], appears to be very similar to PPPoA. However, the PPP protocol is extended to the workstation, and then encapsulated in both Ethernet and ATM. Though this adds a minor degree of complexity to the

connection, it counterbalances that issue by applying all of the benefits of PPP over the entire length of the connection.

REFERENCES

[1] W. Simpson, The Point-to-Point Protocol (PPP), July 1994; www.ietf.org/rfc/rfc1661.txt

[2] L. Mamakos, K. Lidl, J. Evarts, D. Carrel, D. Simone, R. Wheeler, A Method for Transmitting PPP Over Ethernet (PPPoE), February 1999; www.ietf.org/rfc/rfc2516.txt

[3] G Gross, M. Kaycee, A. Lin, A. Malis, J. Stephens, PPP Over AAL5, RFC 2364, July 1998; www.ietf.org/rfc/rfc2364.txt

[4] G. McGregor, The PPP Internet Protocol Control Protocol (IPCP), May 1992; www.ietf.org/rfc/rfc1332.txt

[5] D. Grossman, J. Heinanen, Multiprotocol encapsulation over ATM Adaptation Layer 5, September 1999; www.ietf.org/rfc/rfc2684.txt