

Implementing PPPoE in a Network

By John O'Keefe & Bob Carrick



Making Broadband Manageable: Be Empowered.

© 2004 Fine Point Technologies, Inc. All rights reserved.

The information contained in this document represents the current view of Fine Point Technologies, Inc. on the issues discussed as of the date of publication. Because Fine Point Technologies must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Fine Point Technologies, and Fine Point cannot guarantee the accuracy of any information presented after the date of publication.

This White Paper is for informational purposes only. Fine Point Technologies Makes No Warranties, Express or Implied, In This Document..

Microsoft, Outlook, Windows, and Windows NT are registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.

Other product and company names mentioned herein may be the trademarks of their respective owners.

Fine Point Technologies, Inc. • 139 Centre Street 6th Floor • New York, NY 10013 • USA

TABLE OF CONTENTS

Table of Contents	2
Introduction.....	3
Ease of Use and Seamless Integration	3
What is PPP?.....	4
What is PPPoE?	4
Other Proposals for PPP over ADSL/DSL	8
Dynamic Host Configuration Protocol (DHCP).....	8
PPP over ATM (PPPoA).....	9
History and Status of PPP over Ethernet	10
PPPoE on a Local Network	10
Summary of Benefits	11
Technical Information.....	13
Pseudo-TTY	13
Discovery Phase.....	13
Session Phase	14
Conclusion	15
Regional Locations Worldwide	16

INTRODUCTION

Internet Service Providers (ISPs) are faced with requirements that may conflict with their objectives. In addition there has been a growing debate within the networking industry regarding broadband architecture at customer premises. ISPs must be able to connect multiple hosts at a remote site through the same customer premise access device, but they also want to provide access control and billing functionality in a manner similar to dial-up services using PPP. In many access technologies, the most cost-effective method to attach multiple hosts to the customer premise access device is via Ethernet. It is important to keep the cost of this device as low as possible while requiring little or no configuration on the part of the end-user.

There are several different methods available for a personal computer (PC) to interact with an ADSL/DSL “modem” to access the growing number of high-speed networks and services. Many of these methods require substantial user configuration of the modem and, in some cases, require users to install ATM network interface cards (NICs) in their personal computers, all of which requires substantial user technical knowledge.

By combining two standards, Ethernet and PPP, into PPP over Ethernet (PPPoE), no more knowledge is required of the end-user than is required to set up standard dial-up Internet access. The PPPoE solution uses existing PC hardware and software, existing Ethernet NICs, and existing ADSL/DSL modems. It requires no special configuration or additions to the customer premise modem or ADSL/DSL access network. Further, PPPoE does not alter in any way the collection of PPP protocols vital to ISPs wishing to deliver ADSL/DSL or other broadband services using their existing network model.

Ease of Use and Seamless Integration

While consumers of Internet service want the faster access speed available through ADSL/DSL services, they do not want to deal with the complex installation and configuration that are often required. In short, to achieve rapid acceptance by consumers and telecommuters, broadband services must be easy to use.

In addition to ease of use, it is important that broadband services integrate into the current network infrastructure as seamlessly as possible, with minimal changes to existing equipment or operation systems. Ideally, a broadband technology deployment should fit directly into the existing “operational model” with little disruption to ongoing subscriber services. The more seamless the integration of ADSL/DSL technology into the existing infrastructure of carriers, ISPs and users, the more rapid the acceptance of ADSL/DSL services.

Consider the operation of the current dial-up access infrastructure. A user places a telephone call to establish physical layer connectivity. Using Dial-Up Networking in Windows®, the user makes a connection to the ISP via an analog modem. When the call is connected, a data link session must be established before any user data can be transferred. This is most commonly done today through a PPP session between the user's computer and the ISP's remote access server. PPP authenticates the user, an IP address is dynamically assigned (DHCP), and PPP negotiates various other connectivity parameters. At this point, the user is connected to the service provider (and the Internet).

Service providers can learn several lessons from the current analog dial-up model when providing ADSL/DSL and other broadband consumer services. Most Internet users find it relatively straightforward to establish a new Dial-Up Networking connection in Windows. Service providers offering ADSL/DSL services can take advantage of this by continuing to offer the user a simple Dial-Up Networking setup for connection to broadband services, even though the connection is an “always on” PVC (permanent virtual circuit). Additionally, consumers are accustomed to accessing the Internet through an inexpensive modem requiring minimal configuration, so the simpler and cheaper the ADSL/DSL access device (modem) is,

the more likely consumers will be to buy the service. Finally, ISPs are accustomed to providing consumer Internet access through PPP sessions. Fortunately, PPP can be easily adapted to broadband services with no changes to the existing protocol, enabling ISPs familiar with PPP to offer broadband services such as ADSL/DSL with no need to re-invent the way they do business.

What is PPP?

Point-to-Point protocol, or PPP, is a communications protocol for transmitting information over standard telephone lines. It is a member of the TCP/IP suite of network protocols. PPP is an extension to TCP/IP that provides additional functionality: It can transmit TCP/IP packets over a serial link. TCP/IP by itself cannot be transmitted over a serial link. This makes it unsuitable for WANs (Wide Area Networks), as it is not feasible to extend an Ethernet network over many thousands of miles. Telecommunications companies, however, have offered serial communications links around the globe for many years. To make TCP/IP work over these serial links, it was necessary to create protocols that could transmit TCP/IP packets over serial lines. The two protocols that do this are SLIP (Serial Line Internet Protocol) and PPP. However, PPP has more features than SLIP and has largely replaced it.

In addition, when public telephone system serial links are used, care must be taken to ensure the authenticity of all communications. A router or server receiving a request via PPP in which the origin of the request is not secure would require authentication. PPP incorporates user name and password security to provide the authentication.

Because of its ability to route TCP/IP packets over serial links and its authentication features, PPP is generally used by ISPs to provide dial-up access to the Internet.

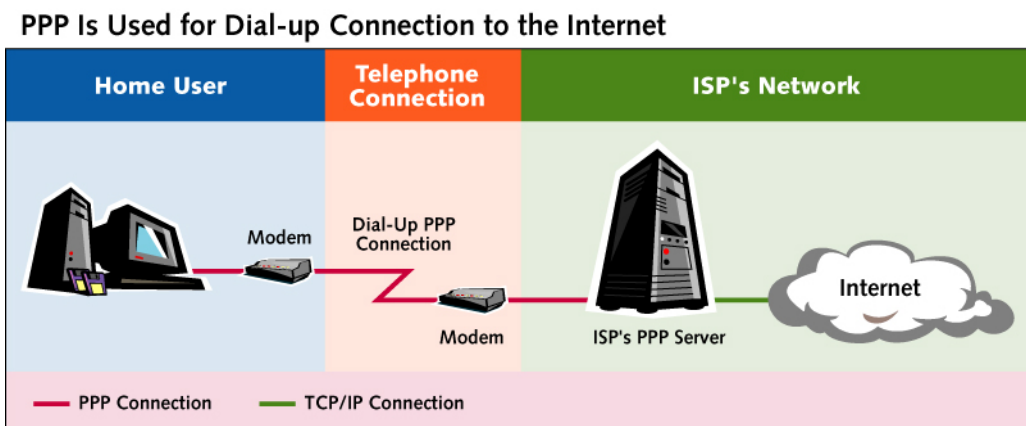


Figure 1

What is PPPoE?

PPP over Ethernet (PPPoE) is PPP (designed for serial communications) that has been adapted to an Ethernet network. Since PPP was designed to do things that are either impossible or unnecessary with Ethernet, there may be some confusion as to why one would *want* to use PPP over Ethernet.

A TCP/IP network and its traffic can be compared to a network of city streets with vehicle traffic. There are many points at which a car can get on or off each street. Additional access points can be added with little disruption. But it is hard to tell how many cars are actually using each street. PPP can be compared to an overhead monorail. Travel is generally between two well-defined points; passengers can only get on and off at those points and need a ticket to board, making it relatively easy to count and monitor

passengers. So PPP over Ethernet is similar to a monorail running over the city street system. It offers speedy access between two well-defined points, and its traffic can be monitored.

PPPoE Allows ISPs to Monitor Internet Traffic Volume

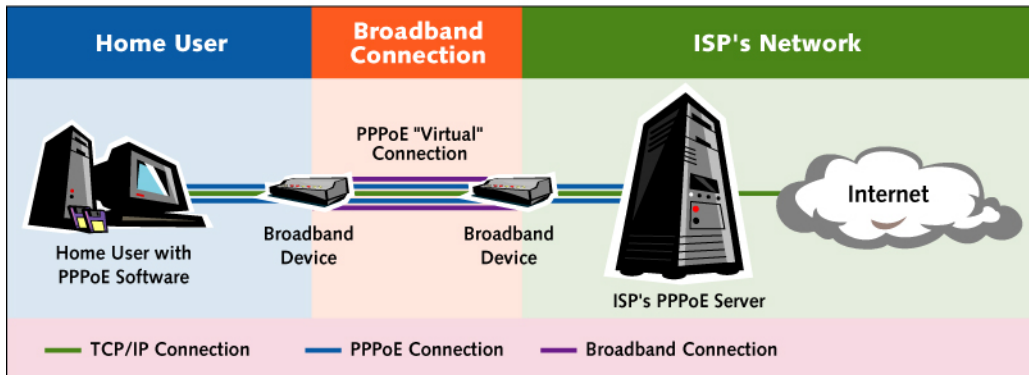


Figure 2

PPPoE provides the ability to connect a network of hosts over a simple bridging access device to a remote access concentrator. Each host utilizes its own PPP stack and the user is presented with a familiar interface. Access control, billing, and type of service can be done on a per-user, rather than a per-site, basis.

To provide a point-to-point connection over Ethernet, each PPP session must learn the Ethernet address of the remote peer, as well as establish a unique session identifier. PPPoE includes a discovery protocol to do this.

PPPoE allows users to utilize industry-standard Ethernet NICs and standard Ethernet drivers to connect to an ADSL/DSL modem. A simple "shim" is used between the existing Windows Dial-Up Networking PPP stack and the Ethernet driver, allowing PPP sessions to take place in standard Ethernet frames. The only requirement of the ADSL/DSL modem is support for Ethernet MAC bridging. The encapsulation on the ADSL/DSL line is standard RFC 1483 Ethernet bridged frames. Multiple users can share an Ethernet segment in the same way they do in corporate networks all around the world today.

PPPoE Uses Ethernet NICs and Drivers

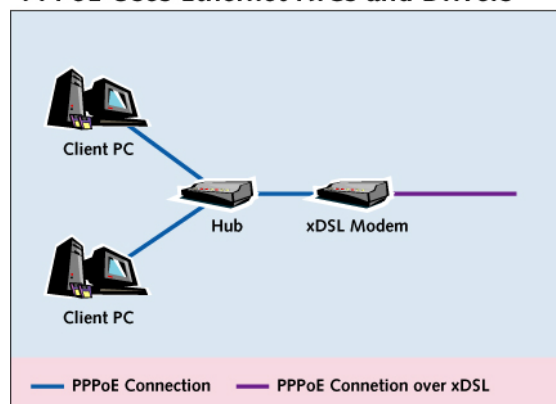


Figure 3

The user takes the following steps to set up ADSL/DSL service using PPPoE:

- Installs a carrier-supplied ADSL/DSL modem pre-configured with a PVC

- Connects the Ethernet port on a NIC in the PC to the Ethernet interface on the ADSL/DSL modem
- Installs the PPPoE driver
- Using Windows Dial-Up Networking, sets up a PPP connection over the Ethernet-connected ADSL/DSL modem
- Clicks on the Dial-Up Networking connection, provides the appropriate user name, domain and password, and clicks “Connect”

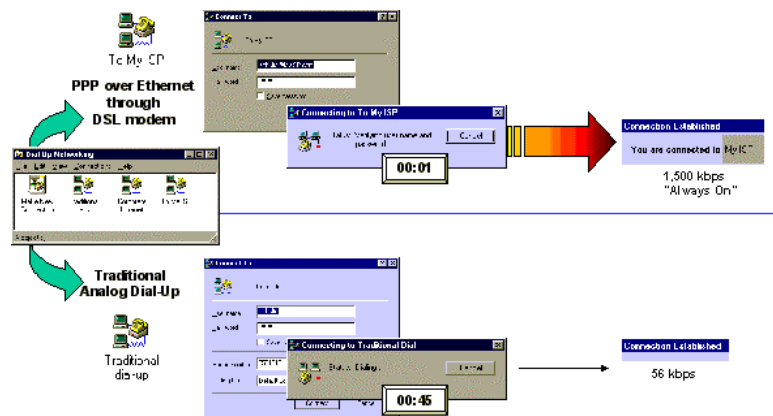


Figure 4: The PPPoE user interface is similar to the PPP user interface

When the user clicks “Connect,” a PPP over Ethernet session is established. This PPPoE session is bridged by the ADSL/DSL modem to an ATM PVC, which connects in an ISP point-of-presence (POP) to a device, such as the ServPoET BMS 1000, capable of terminating an ADSL/DSL PPP session. At this point, the user has established a connection to the ISP using procedures virtually identical to those used to set up a dial-up analog connection but with the faster connection speed and greater bandwidth of ADSL/DSL. Importantly, the entire collection of PPP protocols is unaltered. The Ethernet is simply used to carry PPP messages between the client and the server. At the service provider’s end, the connection session emulates a standard PPP session. Also, additional users can initiate PPPoE sessions using the same ADSL/DSL modem and line, and no additional circuits are required. One circuit can support a number of PPP sessions, simplifying the configuration in the carrier central office.

One PVC Can Support Several PPPoE Sessions

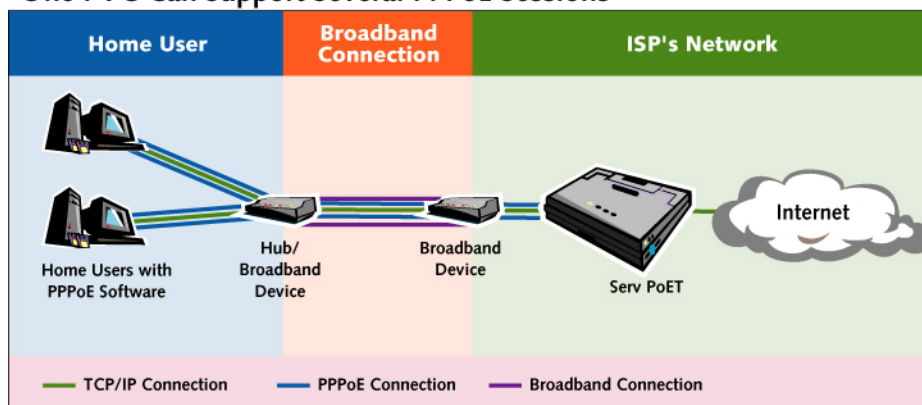


Figure 5

PPPoE offers a solution for providing high-speed, broadband Internet access that simplifies user configuration, utilizes standard low-cost Ethernet NICs, and provides a familiar user interface. In

addition, PPPoE strictly adheres to the PPP standard, works with all existing ADSL/DSL modems, and requires minimal additional driver software (WinPoET client) on the user PC.

OTHER PROPOSALS FOR PPP OVER ADSL/DSL

In selecting the best architecture for their broadband services, service providers want cost-effective service selection and the accounting support required by new and evolving broadband business models. Recognizing the importance of PPP in hastening the acceptance of ADSL/DSL (and other broadband) services, there are currently several proposals for PPP over ATM over ADSL/DSL being promoted within the industry. All of them share a common thread—they are focused on delivering PPP over ATM over ADSL/DSL across the local loop to a carrier central office (CO) then across a regional data network to an ISP POP. They differ only in the architecture that provides communication between the end-user and the ADSL/DSL modem at the customer premises.

Dynamic Host Configuration Protocol (DHCP)

Dynamic Host Control Protocol (DHCP) is one of the alternative protocols available. While DHCP has potential in certain applications, PPPoE meets the broadest range of requirements and enables service providers to utilize current hardware while leveraging investments in legacy provisioning and billing systems.

An architecture based on DHCP offers increased flexibility compared to an architecture based on static IP addresses because it uses DHCP servers that automatically assign IP addresses and configure PCs, providing network access that is transparent to end-users. But, as with static IP addresses, DHCP architecture cannot authenticate end-users (supporting a fee-for-service business model that optimizes broadband revenue streams) unless proprietary and complex software is added.

“Walled garden” is one example of an architecture that adds authentication capabilities to DHCP through add-on software. The broadband user is unknown to the network when first granted access through the DHCP server, so limited access is assigned while the authentication process is completed. In other words, the user is granted access to the network, but it is restricted to a specific area. The surrounding architecture must be able to track all subscribers, determine which are authorized users, and then provide access only to that subset.

Walled garden is extremely complex, requiring proprietary real-time interfaces between the DHCP server, a RADIUS server (for authentication), the Broadband Access Server (BRAS), and a billing server. Difficult to implement, DHCP walled garden also presents maintenance and administrative challenges because so many different applications must be tightly integrated in order to perform the authentication procedure. But even when this approach is working smoothly, it is constrained by the fact that it relies on granting some level of access before the authentication procedure can begin—clearly a potential security problem.

PPP is the most proven architecture, having worked well in the dial-up arena for over a decade. By requiring a password/ID handshaking procedure before access is granted, this approach inherently supports the authentication process required to track usage and bill for service accordingly. The PPP architecture also incorporates the standard RADIUS protocols already at the heart of virtually all customer provisioning and billing systems. As a result, no changes to proven back-end systems are required when adding broadband services. In other words, PPP empowers providers to protect existing investments while creating the new broadband services required to distinguish themselves and capture new customers in competitive markets. At the same time, because PPP supports open access, the market is dramatically increased to encompass the entire population of broadband users—without the need to make new billing infrastructure investments.

PPP over ATM (PPPoA)

The key benefit of PPPoA is its ability to deliver end-to-end Quality of Service (QoS), including guarantees of latency and bandwidth availability. However, this approach requires an ATM connection in the subscriber PC. This adds cost and increases deployment complexity because ATM network interface cards are complex and not always compatible with the desktop operation system (and even when they are, drivers must be configured). In addition, to take full advantage of PPPoA, the network has to support Switched Virtual Circuits (SVCs), which are not widely available. And finally, PPPoA software is not available for all platforms and does not support home LANs or cable or wireless access.

Because of ATM's prevalence in the network core, DSLAMs initially deployed at the network edge were optimized to work with ATM technology. ATM transports voice and data traffic in 53-byte cells. Therefore, ATM-based DSLAMs break up voice and data traffic into ATM-sized cells. Slicing traffic into equal-size cells effectively hides IP datagrams, which contain information such as who generated the traffic, the type of traffic (voice versus data), and where the traffic is going. While this does not cause extensive problems for providing standard high-speed Internet access service, it becomes an increasingly serious disadvantage as DSL providers offer a wider range of concurrent, high-value services for a growing subscriber base. Furthermore, ATM DSLAMs cannot tell one subscriber's traffic from another's. DSL service providers who rely on ATM DSL access technology must have a single, end-to-end PVC from the customer premise DSL router or modem, through the DSLAM, through one or more ATM aggregators or switches, to the Internet router *for each subscriber*. So an ATM-based DSL provider who signs up 100 customers for Internet access must provide 100 end-to-end PVCs. By comparison, PPPoE has none of these drawbacks and even supports QoS features inherent in the network architecture. In fact, the most unfavorable issue mentioned about PPPoE as the ideal architecture for broadband services is that it requires third-party client software.

Despite the negative perception some may have with client software, this is one of PPPoE's strengths because it allows ISPs and CLECs to brand and control their service in a way that would not be possible without it. Since they do not own the infrastructure or control the CPE, the only way these companies can deliver consistent service is through the software they control and provide to customers.

In addition, well-designed, third-party PPPoE client software, unlike the basic PPP drivers bundled with operating systems and used for PPPoA, can provide a range of operational benefits to both the subscriber and the service provider. Chief among these are network management and diagnostic capabilities that can identify operational problems and automatically offer resolutions. This data, available to help desk staff, can also dramatically reduce the time it takes to resolve the problems of customers who call for assistance.

HISTORY AND STATUS OF PPP OVER ETHERNET

PPPoE was submitted in August 1998 to the Internet Engineering Task Force (IETF) as an Internet Draft. Subsequently, a Birds of a Feather (BOF) session was held, and the Internet Draft was updated. In February 1999, rfc2516 was published.

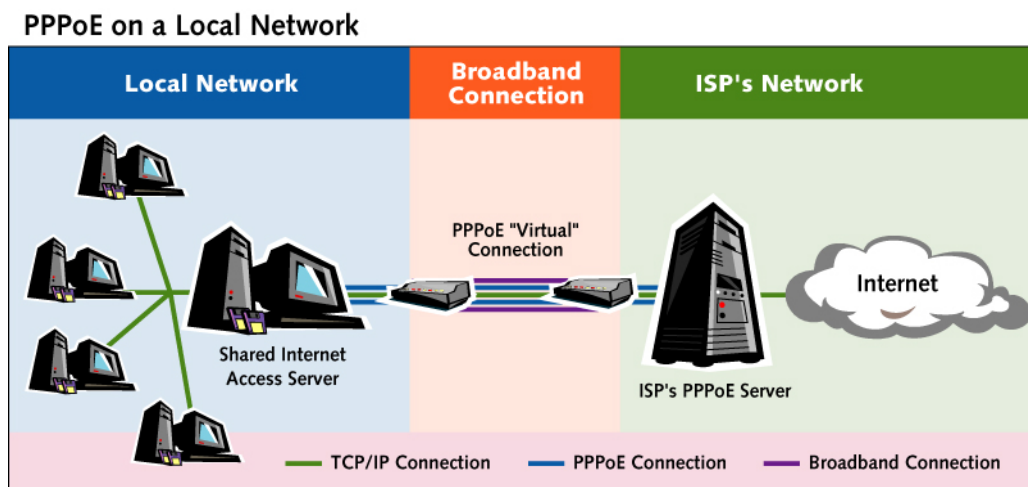
PPPoE on a Local Network

The PPPoE standard requires that PPPoE software place an additional header at the beginning of each TCP/IP packet. This may cause the packet to become larger than the maximum allowable size. Some software solutions handle this transparently, but some require modification of the TCP/IP settings on all of the client computers on the LAN.

Relying on two widely accepted standards, Ethernet and PPP, PPPoE specifies how a host PC interacts with a broadband modem (xDSL, wireless) to achieve access to the growing number of high-speed data networks.

PPPoE requires no more knowledge of the end-user than is required to set up standard dial-up Internet access. Additionally, PPPoE does not require any major changes in the operation model for service providers and carriers.

Figure 6



SUMMARY OF BENEFITS

PPPoE offers many advantages for DSL service providers, as listed below.

Maximize IP address pool

Because PPPoE sessions are the same as PPP sessions, IP addresses are dynamic. Service providers can ensure that users' assigned IP addresses are changed each time they connect.

Because PPPoE emulates a "session" over Ethernet, service providers can charge based on connect time. This allows them to discourage permanent connections and to over-subscribe their IP address pool.

Bill based on user not location

Because PPP sessions almost always require authentication, DSL service providers can bill the correct client regardless of connection location (as dial-up ISP's can now).

Use other protocols

PPPoE can encapsulate non-IP protocols. Any protocol that can be encapsulated by PPP can be sent via PPPoE.

Offer user authentication as a value-added service

Service providers can enter into agreements with large organizations to authenticate users and provide dedicated sessions behind the organizations' firewall (for employees who need remote access, for example). RADIUS requires that the user be authenticated before establishing a connection to the Internet. With PPPoE (via WinPoET/MacPoET), a user must log in via WinPoET/MacPoET software, which, in turn, authenticates them for the RADIUS server.

Increase speed of broadband deployment

Service providers can leverage their existing infrastructures with broadband technologies. The primary component of PPPoE infrastructure is RADIUS authentication; and typically, service providers have already implemented RADIUS authentication for their dial-up service.

Reduce administration costs

Service providers can provision or un-provision broadband accounts using the same methods as for dial-up accounts, utilizing existing billing management software already integrated with RADIUS. A subscriber's account can be simply deactivated through the billing software by removing the RADIUS login, and the subscriber will no longer be able to connect to the Internet. The subscriber can also be re-activated through the billing system.

Prevent theft of service

A smart user can purchase a hub and set up a home network with a residential ADSL account, thus accessing one account through more than one computer. PPPoE (via WinPoET/MacPoET) requires that

the user log in. With RADIUS, the number of simultaneous logins for a user can be restricted, thereby limiting the number of computers that can be logged in at one time.

Provide value-added services

Many third-party products require the use of a RADIUS database to manage user settings. These products allow a service provider to compete more effectively by offering such services as Web and e-mail filtering.

Allow for broadband resellers

With PPPoE, agent resellers can resell a broadband provider's service. The agent resellers act as Internet Service Providers and can control and manage their own users by implementing their own RADIUS servers.

New service offerings and revenue sources for resellers

Service Providers now have the option of leveraging PPPoE to build customized services that create new revenue streams on top of their existing IP "transport" offerings. Service Providers who have limited their Internet access offerings to leased lines and dial-up services can use PPPoE to expand their service to include the DSL/ADSL market space. Multiple PCs at a customer site – be that a small business, and multiple unit dwelling or a private home with two or more PCs – can all have discrete connections over a shared line. Each to different service providers each billed separately. Further – a single PC can accommodate multiple services, again, over a single line, and with each service producing a new revenue stream.

TECHNICAL INFORMATION

PPP over Ethernet uses two new Ethernet type codes: one for PPP control messages and one for IPCP data. The frame formats are completely unchanged from Ethernet V2 as outlined below:

DA	SA	PPP Control	Data	FCS
DA	SA	IPCP Data	Data	FCS

The client requires a "shim" between the existing Dial-Up Networking PPP stack and the NDIS Ethernet driver, allowing Windows to bind a PPP session to an Ethernet adapter.

The protocol works as follows: The user begins a PPP session by performing the usual, familiar steps (using Dial-Up Networking in Windows). The PC sends an LCP config request message in a broadcast Ethernet frame using the PPP control type code. Any standard MAC layer bridging modem forwards this frame. The device terminating the PPP session (e.g., a router or a ServPoET device) responds to the PPP message with the normal LCP response by unicasting the message back to the client PC using the same frame format. ServPoET maintains the binding between the local MAC address and the PPP session ID. LCP, PAP/CHAP, and IPCP control messages are unicasted between the two endpoints in this manner until the PPP state machine reaches the normal data transfer stage. Data is then transferred using the IPCP data Ethernet type code. All PPP control and data packets are raw Ethernet encapsulated.

Also, if additional user PCs initiate PPP sessions using the same ADSL/DSL modem and line, no PVCs are required. One connection can support a number of PPP sessions because the sessions are demultiplexed by MAC address, and the remote encapsulation is RFC 1483 bridged, not VC multiplexed.

ServPoET is based on the Linux operating system for the following key reasons:

Pseudo-TTY

The pppd program and the Linux kernel expect to transmit PPP frames over a TTY device. UNIX (and Linux) supports the concept of a pseudo-TTY. This is a device that emulates a TTY, but instead of being connected to a physical terminal, it is connected to a UNIX process. When something writes to the pseudo-TTY, the data appears on the standard input of the back-end process. When the back-end process writes to its standard output, the data can be read from the pseudo-TTY.

Pppd also supports a pty option. This option automatically starts the back-end process and performs all the mundane operations required to connect it to a pseudo-TTY. So the PPPoE link is started by starting pppd with the appropriate pty option, which runs the PPPoE executable connected to the pseudo-TTY.

Discovery Phase

When PPPoE begins executing, it starts PPPoE discovery. It creates a raw Ethernet socket. This special socket allows user-space programs to transmit and receive raw Ethernet frames. PPPoE constructs and transmits a PADI frame and waits for PADO frames. When a PADO frame arrives (if it meets the criteria specified on the PPPoE command line), PPPoE transmits a PADR frame. After it receives a PADS frame, it records the session ID and moves to the session phase.

Filtering only on the Ethernet frame type is acceptable, and Linux raw sockets can perform this level of filtering. (Non-PPPoE frames in the kernel need to be filtered out; otherwise, non-PPPoE traffic could consume large amounts of CPU time as PPPoE is scheduled in to read the frame.)

Session Phase

When the session phase begins, PPPoE reads asynchronously framed PPP data on standard input and writes it to standard output. When a frame is transmitted over the PPP interface, PPPoE's standard input becomes readable. PPPoE reads from standard input and collects data until it has assembled an entire PPP frame. PPPoE must keep a small state machine to record where it is in the PPP frame assembly.

During PPP frame assembly, PPPoE removes escape sequences and “de-stuffs” the frame. This converts the asynchronous PPP framing into a synchronous PPP frame. After the PPP frame is assembled, PPPoE headers are added, and the frame is transmitted over the raw socket. (Most of the PPPoE headers are constant for a given session, so the PPP frame is simply assembled into a buffer right after the constant PPPoE header portions.)

When the Ethernet card receives an incoming PPPoE frame, PPPoE's raw socket becomes readable. In this case, a read operation will return one (and only one) frame and will return the entire frame if the buffer is big enough. PPPoE reads the frame into a buffer. It then adds asynchronous byte stuffing to the data and computes the PPP frame-check sequence (the PPP FCS is not transmitted over PPPoE). Finally, it writes the result to standard output where it is picked up by the kernel or pppd.

CONCLUSION

While there are several valid approaches to delivering PPP over ADSL/DSL and other broadband services, some may be difficult for the end-user to install, configure, and manage. By creating end-user complexity, these methods may in turn hinder widespread deployment and adoption of consumer ADSL/DSL and other broadband services. This paper outlines an alternate approach known as PPP over Ethernet (PPPoE) that dramatically simplifies the process of using PPP to deliver broadband services. PPPoE uses industry-standard, low-cost Ethernet hardware, standard Windows Dial-Up Networking, and works with all existing ADSL/DSL modems.

REGIONAL LOCATIONS WORLDWIDE

United States (Corporate Headquarters)

139 Centre Street, 6th Floor
New York, NY 10013
USA
+1.212.962.7410
info@finepoint.com

United Kingdom

Globix House
1 Olivers Yard
London EC1 Y1HQ
UK
+ 44.2075.264818

France (Europe, Middle East, Africa)

Les Algorithmes, Bât. Aristote A
2000, Route des Lucioles, BP 29
06901 Sophia Antipolis
France
+ 33.1707.18418