On the Establishment of an Access VPN in Broadband Access Networks

Reuven Cohen, Technion

ABSTRACT

Many corporate networks use the Internet as a cost-effective substitute for expensive leased lines and long-distance telephone calls, for allowing their users to have on-demand connectivity into their intranets through ad hoc tunnels, a concept known as access VPN. Establishing an access VPN in a broadband access network is very often a difficult networking challenge that requires several tunnels to be established between the various involved entities. This is especially the case in an open broadband access network where the association between a host and its ISP is not known to the network in advance. This article discusses several schemes for establishing an access VPN in a broadband access network, and explains the need for the various tunnels in each scheme.

INTRODUCTION

As the Internet becomes a popular low-cost backbone infrastructure, its universal reach has led many companies to consider constructing a secure virtual private network (VPN) over the public Internet. Essentially, a VPN is a private data network that uses a nonprivate data network to carry its traffic. VPNs offer an alternative to the traditional leased line or frame relay networks by utilizing an established public network. The most ubiquitous, least expensive nonprivate data network is the Internet, which is the perfect foundation for a VPN. The challenge in designing a VPN is often to provide the security and address flexibility of the traditional private self-administered corporate intranet over the nonprivate backbone.

There are several possible VPN applications. The most popular are the access VPN and intranet/extranet-VPN. An access VPN allows remote corporate users to have on-demand connectivity into their corporate Intranets through ad hoc tunnels. An intranet/extranet VPN links the network of a headquarters office to the networks of remote branches (intranet) and potentially to the networks of business partners such as vendors, providers, or distributors (extranets).

There is no single definition of an access VPN. However, in this article we define an access VPN as a scheme that allows secure remote access to an internal corporate server. Such a scheme should fulfill the following requirements:

(R1) User authentication and authorization: The scheme should be able to identify the user and to verify that this user is authorized to access the contacted internal server.

(R2) Data privacy: The scheme should guarantee that the exchanged data is encrypted and authenticated at least when it is sent over the public Internet.

(R3) Private addressing: Many corporations use private IP addresses in their intranets. In these cases the access VPN scheme should usually be able to assign the remote user a private IP address taken from the same range. This is not a trivial requirement if the packet has to pass through the public Internet.

The most important mechanism used by VPNs is the concept of tunneling. The idea behind this concept is that a part of the route between the originator and the target of the packet is determined independent of the destination IP address. The importance of tunneling in the context of access VPNs in broadband access networks is twofold. First, the destination address field of a packet sent in an access VPN may indicate a non-globally-unique IP address of a corporate internal server. Such an address must not be exposed to the Internet routers because these routers do not know how to route such packets. Second, very often a packet sent by a user of an access VPN should be forwarded first to the ISP of this user, and only then from the ISP toward the corporate network. In such a case the first leg of the routing — between the host and the ISP — cannot be performed based on the destination IP address of the packet, even if this address is globally unique.

The main reason for the wide variety of access VPN solutions and for their complexity is that up to five entities can be actively involved as tunnel endpoints: the end host (user's PC), the broadband modem, the operator access gateway, the ISP access gateway, and

0163-6804/03/$17.00 © 2003 IEEE

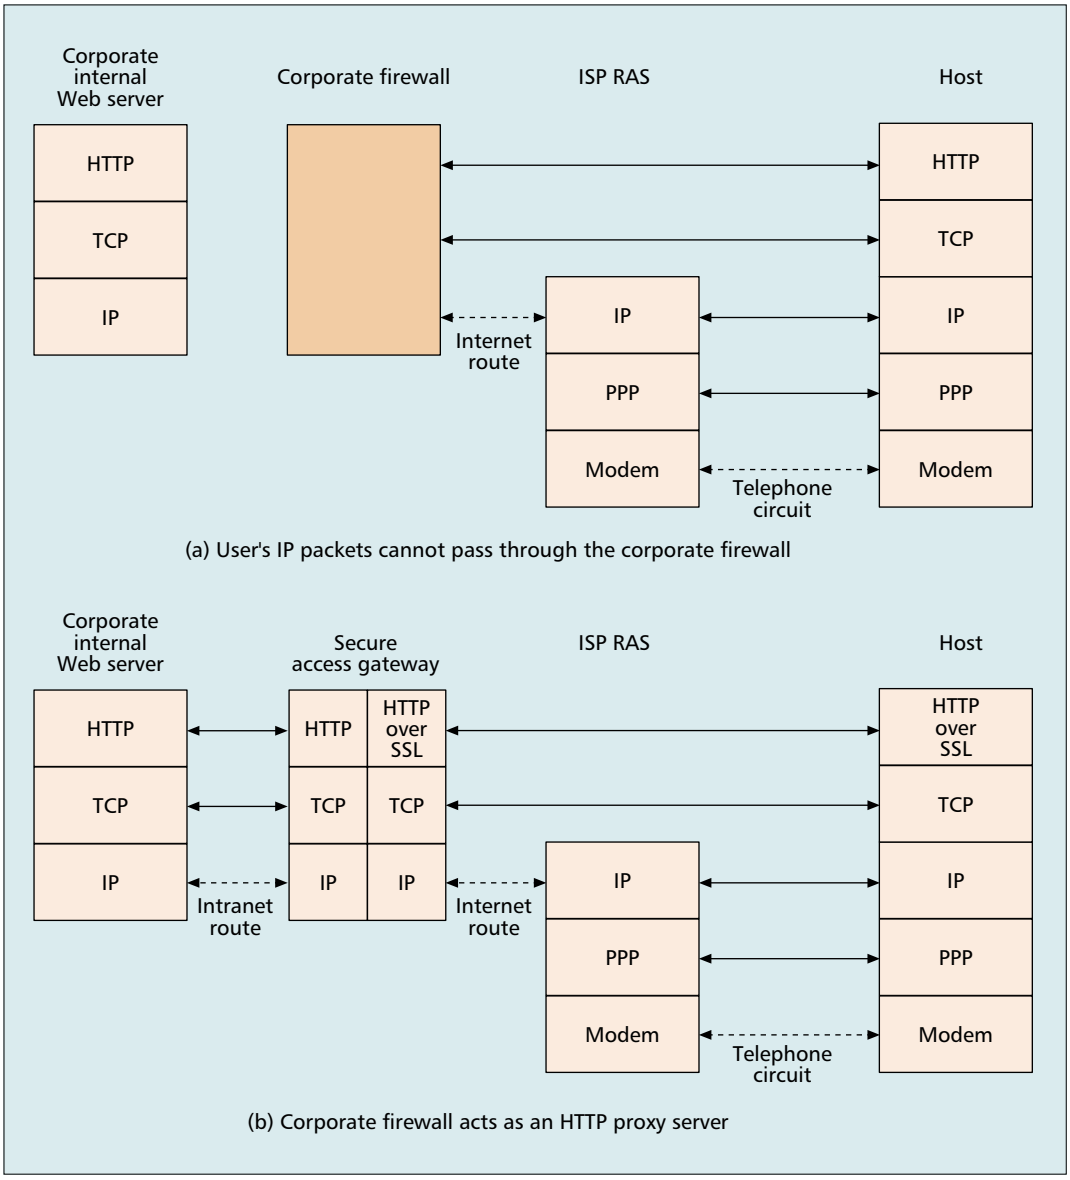

Figure 1. *PSTN dialup to the ISP RAS (scheme 1).*

PPP's LCP allows the connecting user to provide authentication information, like a user-name and a password. PPP's NCP allows the ISP to configure the connecting host with networking parameters, like the IP address assigned to the host and the IP address of the ISP's DNS server.

the corporate access gateway. The purpose of this article is to discuss the various schemes for establishing an access VPN in a broadband access network, and in particular to explain the need for the various tunnels required in each scheme.

The rest of the article is organized as follows. We start by discussing the concept of the access VPN in the traditional dialup access environment. This background is needed because many access VPN solutions developed in the context of broadband access networks try to emulate the dialup access environment in order to streamline the transition of all the involved entities from the narrowband access world to the broadband access world. We then discuss the generic structure of a broadband access network, and explain the role each of the involved entities plays in the context of an access VPN. We then present the various approaches for establishing an access VPN in a broadband access network. Finally, we conclude the article.

Establishing an Access VPN in a Dialup Access Network

In this section we address the relatively simple case of dial access, where the user employs an analog modem and dials into a bank of modems called a remote access server (RAS). There are two basic modes for a dialup access VPN. In the following discussion these modes are referred to as schemes 1 and 2. In scheme 1, the user's modem connects to the Internet through an Internet service provider's (ISP's) RAS. The protocol stack for this scheme is depicted in Fig. 1a. The most important component in this scheme is the Point-to-Point Protocol (PPP) [1], executed over a public switched telephone network (PSTN) circuit between the user's analog modem and the ISP's RAS. PPP can be viewed as a method for encapsulating an IP packet over a serial link. However, it has two additional important components [2]:

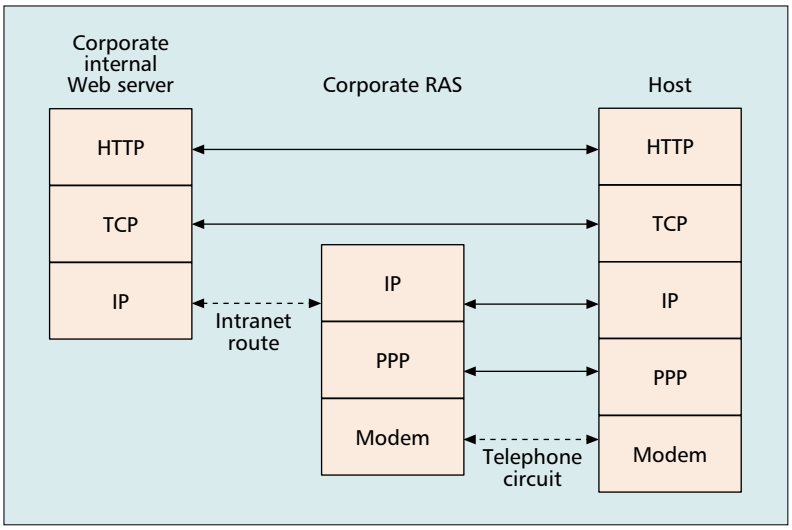

Figure 2. *Direct dialup to the corporate RAS (scheme 2).*

- A link control protocol (LCP)
- A family of network control protocols (NCPs) for establishing and configuring different layer 3 protocols. The network control protocol for IP is called IPCP.

PPP's LCP allows the connecting user to provide authentication information, like a username and a password. PPP's NCP allows the ISP to configure the connecting host with networking parameters, like the IP address assigned to the host and the IP address of the ISP's DNS server. The PPP link between the end user and the ISP can be secured using encryption [2].

After being connected to the ISP, the user host can send IP packets to any server connected to the public Internet. However, corporate intranet servers cannot be accessed by such a user because requirements (R1)–(R3) are not fulfilled for the following reasons:

- The user is authenticated by the ISP, but not by the corporate servers.
- PPP encryption, when used,[1] secures only the packets traversed between the ISP and the user over the relatively secure PSTN circuit, but not the packets traversed between the ISP and the corporate intranet.
- The IP packets sent by the user host carry in their source IP address field the IP address assigned to this host by the ISP, not a legal corporate Intranet's IP address.

Due to these constraints, the corporate firewall will probably be configured to block these packets from entering the corporate intranet.

One way to address these problems is to employ a special application layer gateway that works as an HTTP proxy server (Fig. 1b). In order to access the intranet servers, the remote user first accesses the corporate proxy server. Using "HTTP secure," namely HTTP over transport layer security (TLS) [3], the proxy server authenticates the remote user. After the user is authenticated, the user's browser continues using HTTP secure with the corporate proxy server. The proxy server authorizes each request, and sends it through the intranet to the proper internal server. Therefore, (R1) and (R2) are fulfilled. Since the IP packets sent by the proxy in the intranet carry the IP address of the proxy in their source address field, (R3) is fulfilled for the client IP address. In order to allow the accessed Web server to have a non-globally-unique IP address, thereby fulfilling (R3) for this address as well, a URL mapping approach can be used. The idea is that the URL will be associated with the globally unique address of the gateway, but it will also contain the name of the internal server. When the gateway receives the request, it translates it into a new request to the internal Web server. The main disadvantages of this approach are that the secure access gateway might become a bottleneck, and that with this approach the user can only work with an application layer protocol that runs over TLS.

The second basic approach for a PSTN dialup VPN is scheme 2, presented in Fig. 2. In this scheme, the corporate intranet serves as the ISP of the dialup user. The user dials up directly to the corporate RAS rather than to the ISP's RAS. The user is authenticated by the corporate RAS, and is provided with an IP address from the corporate pool of addresses. This solution fulfills (R1)–(R3), but has two cost-related drawbacks:

- RAS cost: The corporate network must deploy and operate an RAS with a pool of modems.
- Dialup cost: Whereas in scheme 1 the user can dial into a local ISP, thus incurring the cost of a local phone call only, in scheme 2 an expensive long-distance phone call may be required. In other words, this scheme does not use the Internet as part of the access VPN, and therefore does not fulfill our definition of an access VPN.

The Point to Point Tunneling Protocol (PPTP) [4] and its successor, the Layer 2 Tunneling Protocol[2] (L2TP) [5], can eliminate these drawbacks. PPTP and L2TP allow two end nodes to emulate a PSTN circuit on an IP network. This circuit is then used for setting up a PPP connection between the two nodes. PPTP and L2TP have two modes: the voluntary tunneling mode and the compulsory tunneling mode. Scheme 3, depicted in Fig. 3a, uses the voluntary tunneling mode. In this scheme, the user dials up to a local ISP. Over the established PPP connection, the user's host is now able to connect to any server in the Internet. The host invokes the PPTP or L2TP protocol in order to set up a PPTP/L2TP tunnel to the corporate network access server (NAS)[3]. Over the established PPTP/L2TP tunnel, the host sets up *another* PPP connection, but this time with the corporate NAS rather than with the ISP's RAS. From the perspective of the host's networking stack, scheme 3 can be viewed as a combination of schemes 1 and 2. The PPP frame sent by the host to the ISP contains two PPP and two IP headers. Moreover, the user's host is assigned two different IP addresses: an IP address for the "lower IP layer" (IP-1) is assigned by the ISP, and is used when the packet is routed over the Internet. The IP address for the "upper IP layer" (IP-2) is assigned by the intranet NAS, and is used when the packet is routed in the intranet.

[1] Because a telephone circuit is considered secure, most PPP over PSTN stacks did not implement the PPP authentication and encryption mechanisms. However, when PPP is tunneled over PPTP, as discussed later, in order to be used in the nonsecure packet-switched network, these mechanisms are implemented and are often considered PPTP security rather than PPP security.

[2] The Layer 2 Tunneling Protocol (L2TP) was developed as a replacement for PPTP and another tunneling protocol called L2F. However, PPTP is still more widespread in the client stack, and is more likely to be used in the near future for voluntary tunneling. L2TP is more likely to be used for compulsory tunneling.

[3] Many papers use the terms RAS and NAS interchangeably. However, we use RAS for a server that receives PPP connections over PSTN circuits, and NAS for a server that receives PPP connections over an IP network.

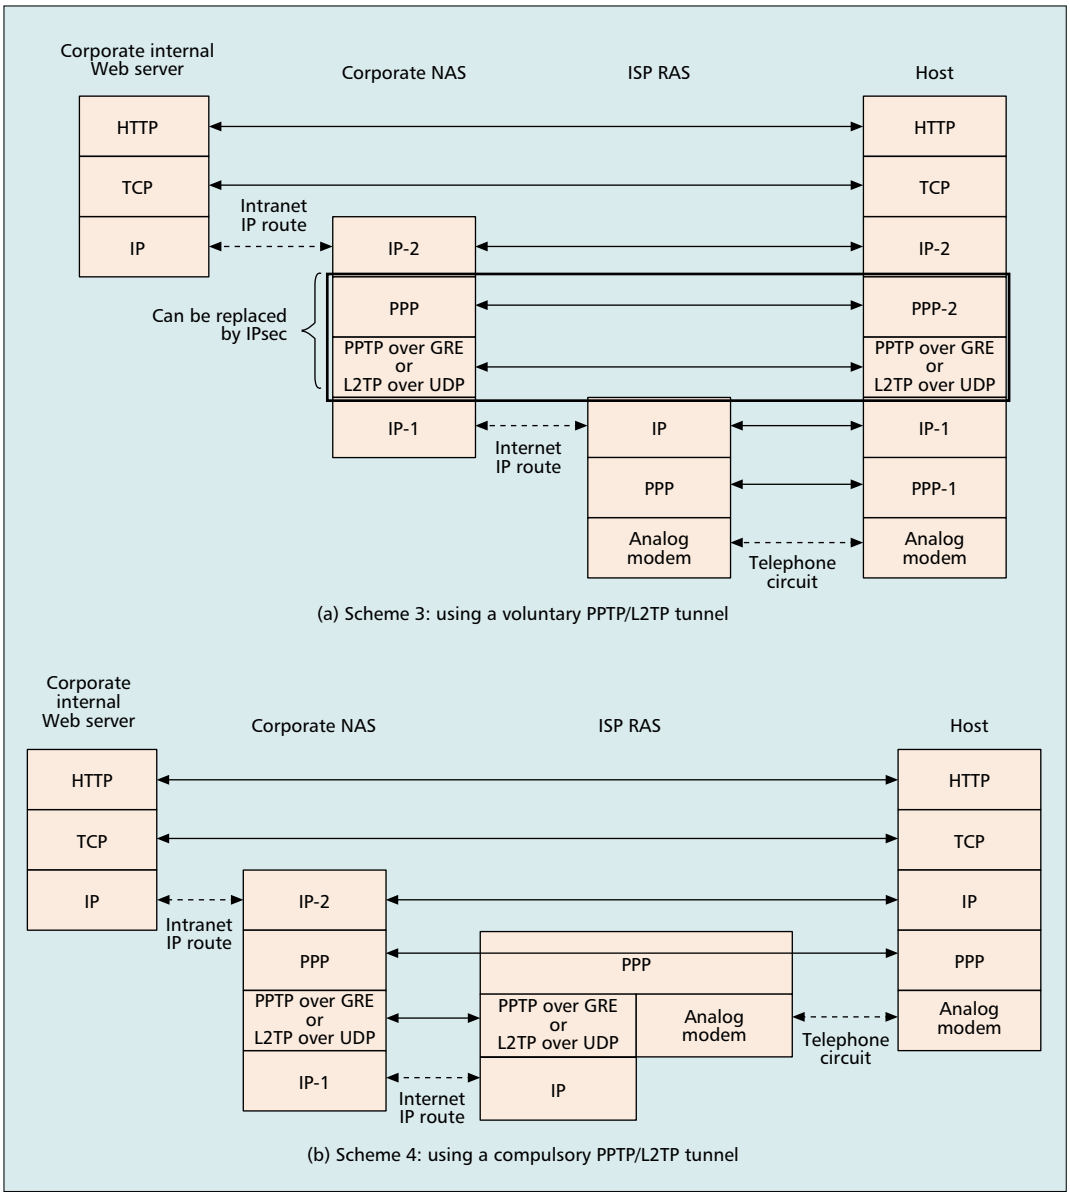

Figure 3. *An access VPN using a PPTP/L2TP tunnel (schemes 3 and 4).*

IPsec is a set of protocols developed by the Internet Engineering Task Force (IETF) to support secure exchange of packets at the IP layer. Typically, IPsec supports an intranet/extranet VPN by creating security associations between gateways at the edge of customer networks. However, IPsec can also be used to support an access VPN by replacing the PPP and the PPTP tunnel between the host and the corporate NAS (Fig. 3a). PPP over PPTP is easier to deploy than IPsec, but provides weaker security. Since the decision whether to use PPTP or IPsec does not affect the number of tunnels and entities participating in each tunnel, in the rest of this article we mention both options when applicable.

The other possible tunneling mode of PPTP and L2TP, compulsory tunneling, is used as follows. The user employs a standard dialup stack and sets up a PPP connection to the ISP. The ISP recognizes this user as a member of a certain intranet. It therefore tunnels the PPP connection over a PPTP/L2TP tunnel to the corporate intranet NAS. This scheme is referred to as scheme 4 (Fig. 3b).

In scheme 3 the PPTP/L2TP tunnel is called a voluntary tunnel because it is set up by the host. In contrast, in scheme 4 the PPTP/L2TP tunnel is set up between the ISP and the corporate NAS with no control of the end host. This tunnel is therefore considered compulsory from the remote user point of view. The main advantage of scheme 3 over scheme 4 is that the ISP does not play any role in the creation of the VPN. The main advantage of scheme 4 is that it is transparent to the end host's networking stack.

In the rest of this article we shall only mention the voluntary tunneling mode of PPTP/L2TP. However, for each scheme that employs this mode, an equivalent scheme that employs the compulsory tunneling mode of PPTP/L2TP can be drawn using the guidelines in Fig. 3.

The main advantage of scheme-3 over scheme-4 is that the ISP does not play any role in the creation of the VPN. The main advantage of scheme-4 is that it is transparent to the end host's networking stack.

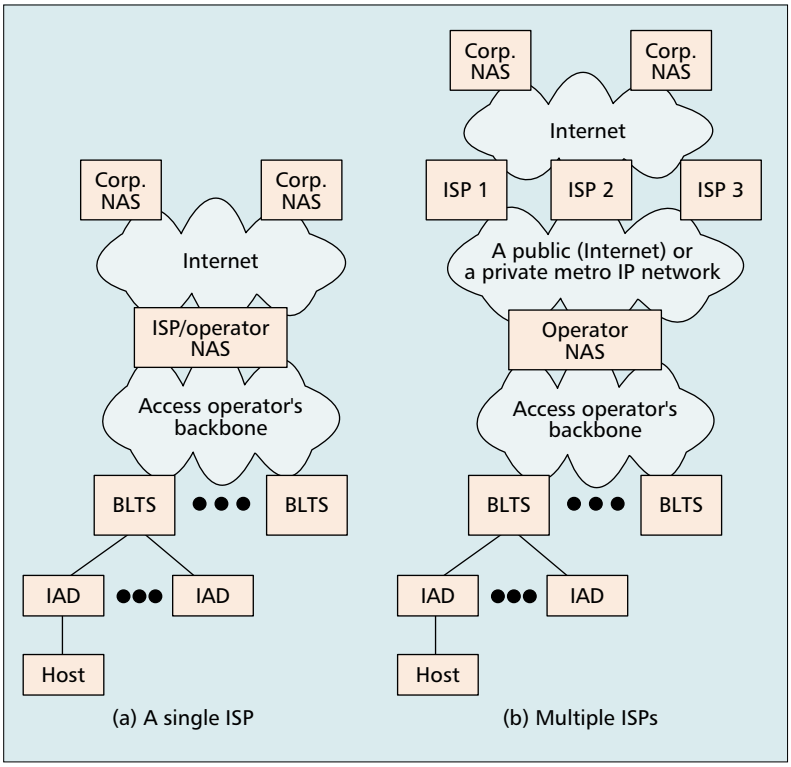

Figure 4. *The structure of a typical broadband access network: a) a single ISP; b) multiple ISPs.*

The Structure of an Open Access Broadband Access Network

Several access technologies are being deployed to deliver broadband services to home users. The key technologies are as follows:

- The digital subscriber line (xDSL) family of technologies that uses the telephony cooper wires as a physical channel.
- Cable modem technology that uses the cable TV infrastructure.
- The wireless family of technologies that uses radio communications instead of terrestrial wires.
- Passive optical network (PON) technology that delivers high-speed rates to business customers over an optical access network. This technology is considered cost effective because it allows fiber sharing and laser sharing among many users.

Despite the different technologies, we can consider a generic access network structure, depicted in Fig. 4. Such a network consists of an integrated access device (IAD) in the user premises that implements one of the above mentioned access technologies, an access physical link, and a broadband link termination system (BLTS) on the operator side of the access link. BLTS is our generic term for the network side of the broadband physical link, sometimes also known as the head-end. It is located in either the operator central office or a street cabinet.

The exact functionality of the access operator's NAS, as well as many other technical issues related to the context of this article, depends to a large extent on a service provisioning policy controversy, often known as "the open access debate." The debated issue is whether the access operator (a cable company, PSTN company, etc.) is allowed to serve as the only ISP of the end host, or must serve as a logical link between the users and their ISPs. These two approaches will be referred to as "the single ISP case" and "the multiple ISPs case," respectively. The case where the access operator is allowed to be one of many ISPs from which a user may select falls into the multiple ISPs case. The two cases are shown in Fig. 4.

At first glance it seems that the multiple ISPs case imposes no new technical challenge, because this is a common practice in the dialup world. However, one of the most important differences between dialup access and broadband access is that in a dialup (PSTN) network a user can establish layer 1 connectivity, by means of a PSTN circuit, with every entity connected to the PSTN network. In contrast, in a broadband access network the end user has layer 1 connectivity and layer 2 connectivity only with the access operator.

Establishing an Access VPN in a Broadband Access Network

We have seen that there are many ways to set up an access VPN over a PSTN-based access network. However, the transition from a PSTN-based access solution to a broadband-based access solution is not straightforward for the following two reasons:

- In most of the schemes related to PSTN access, only three parties are involved: the remote user, the ISP's RAS/NAS, and the corporate RAS/NAS. However, in most broadband access schemes the access operator also plays a key role (Fig. 4a vs. Fig. 4b).
- As shown earlier, all the schemes for dialup access VPNs employ a PPP-based stack at the host. However, hosts in a broadband access network usually use an Ethernet rather than a PPP networking stack, because the connectivity between the host and the broadband modem is Ethernet-based. Therefore, running a PPP in such a stack requires tunneling.

The construction of an access VPN in a broadband access network depends to a large extent on the approach by which the remote host is connected to its ISP. Since connectivity of the host to its ISP in the single ISP case is different from the multiple ISPs case, we address these two cases separately in the following discussion.

The Single ISP Case

We saw earlier that PPP plays a key role in the setup of a PSTN-based access VPN. However, in most cases a host connected to a broadband access network must have an Ethernet-based rather than a PPP-based networking stack. A possible way to get networking configuration parameters from the ISP in the single ISP case

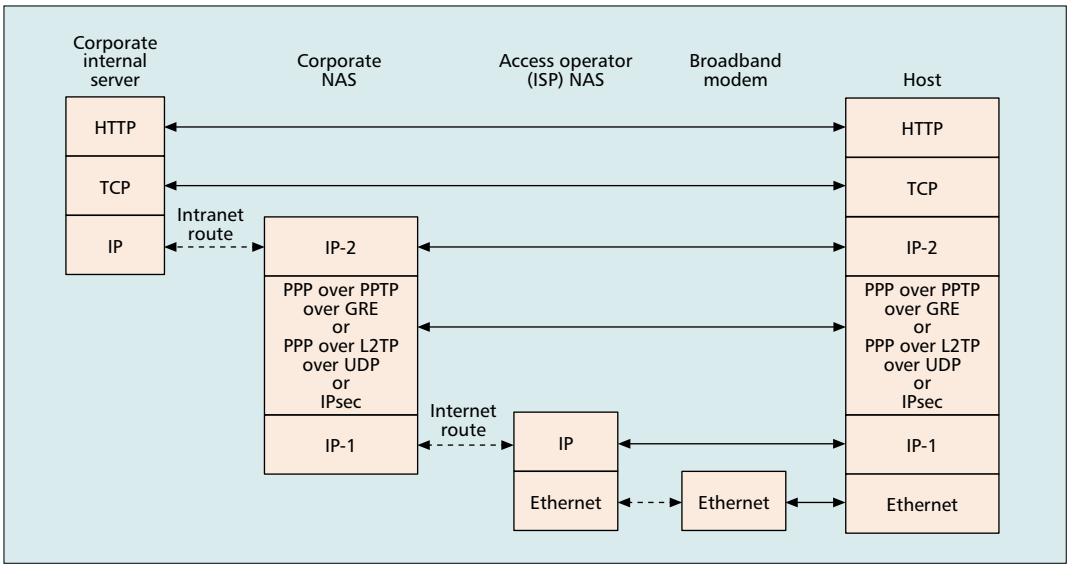

Figure 5. *An IP-based access to the ISP (scheme 5).*

without using PPP's IPCP is using the Dynamic Host Configuration Protocol (DHCP) [6]. DHCP requires the host to have layer 2 connectivity with the access operator's DHCP server or a DHCP relay server. Authentication and authorization can be performed either by the DHCP server, based on the host's MAC address, or after the host gets its initial configuration using HTTP-based tools. After getting connected to the ISP, the user sets up a PPTP/L2TP voluntary tunnel, or an IPsec tunnel, with the corporate NAS as shown in Fig. 5 (scheme 5).

In the description above it was assumed that the single ISP is also the access operator. When this is not the case, the ISP can pre-allocate the DHCP server of the operator a pool of IP addresses for the local hosts. Another option is that the DHCP server of the operator will act as a DHCP relay that forwards the DHCP messages issued by the local hosts to the DHCP server of the ISP and vice versa.

The Multiple ISPs Case

In the multiple ISPs case each host in the broadband access network may have a different ISP. Two issues in this case are as follows:

- How does a user inform the operator of the selected ISP?
- How is the user's host configured with network parameters of the selected ISP?

We shall further distinguish between two subcases: multiple ISPs with static service selection and multiple ISPs with dynamic service selection. In the first subcase each host is pre-associated, using some offline mechanism, with an ISP. Hence, the first issue does not exist. In the second subcase the user is allowed to select an ISP and switch from one ISP to another using some online mechanism.

Scheme 5, as described above for the single ISP case, is still applicable for multiple ISPs with static service selection. When the DHCP server of the operator receives the DHCP message of the host, it knows the ISP with which the host is associated, and processes this message accordingly, by either allocating a set of network configuration parameters for this ISP or relaying the DHCP message to the DHCP server of this ISP.

When dynamic service selection is required, scheme 5 has a chicken-and-egg problem because the host receives network parameters using DHCP before the user is able to execute some application layer protocol, like HTTP, in order to select an ISP. One way to address this problem is by employing the concept of network address translation (NAT) [7]. The idea is that the host gets a set of networking parameters, and in particular an IP address, from the operator's DHCP server. The IP address can be used only for contacting the service portal of the operator. The user uses HTTP in order to select an ISP from the service portal. Consequently, the operator's NAS contacts, usually using the RADIUS protocol, the selected ISP, and gets a new IP address for the user's host. Next, each IP packet sent by the host not to the operator's portal is processed by the NAT logic of the operator, and the source IP address is replaced with the one assigned by the ISP.

However, ISPs and access operators tend to avoid NAT because it has troubles with some application layer protocols. In order to dispense with NAT while still employing an IP-based rather than a PPP-based access scheme, a new extension for DHCP can be used. This new extension is referred to as *force renew* [8]. The idea is that after the user uses HTTP in order to select an ISP from the service portal, the operator's DHCP server sends a unicast FORCERENEW message to the host. Upon receipt of this message, the host changes its DHCP state to RENEW, and tries to renew its lease according to the normal DHCP procedure. The DHCP server replies to the DHCP REQUEST with a DHCP NAK. Consequently, the host must broadcast a new DHCP DISCOVER message to start a new handshake with the DHCP server during which it is assigned a new set of network configuration parameters.

Another solution for multiple ISPs with dynamic service selection is to adapt PPP to an

When "dynamic service selection" is required, scheme-5 has a chicken-and-egg problem because the host receives network parameters using DHCP before the user is able to execute some application layer protocol, like HTTP, in order to select an ISP.

Like L2TP and PPTP, PPPoE is also considered a Layer-2 tunneling protocol, because it allows a Layer-2 protocol (PPP) to be tunneled over another layer. In fact, at first glance it seems that PPTP/L2TP can replace PPPoE for the tunneling of the PPP connection to the ISP.

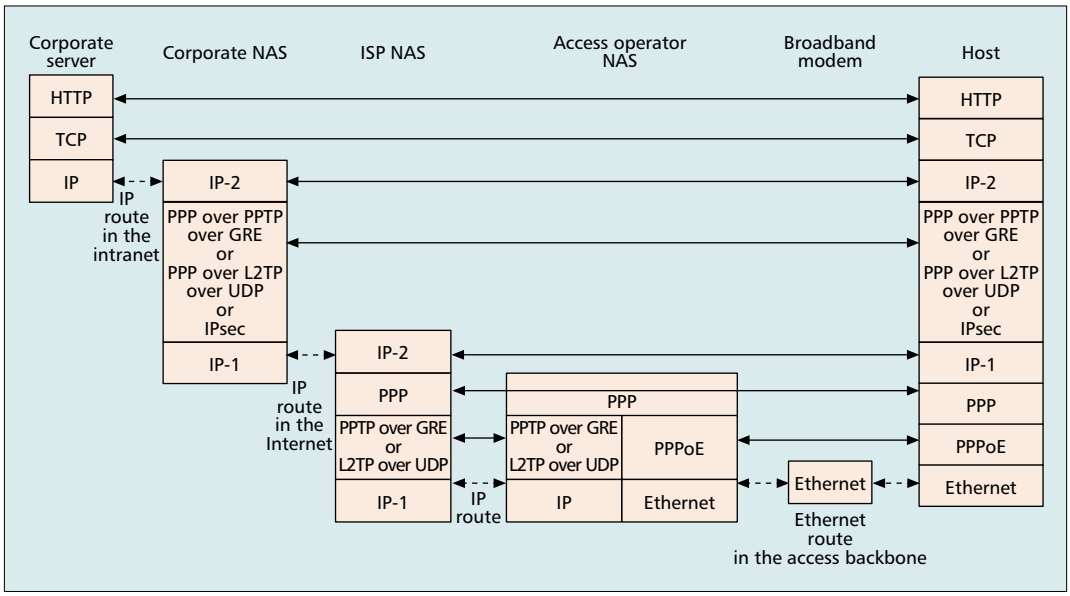

Figure 6. *PPP is tunneled to the access operator over PPPoE, and then to the ISP over a compulsory L2TP/PPTP tunnel (scheme 6).*

Ethernet-based networking stack. To this end, a new layer, called PPP over Ethernet (PPPoE), was developed for the host protocol stack [9]. This layer allows PPP to be tunneled over an Ethernet route that emulates the role of a phone circuit. The Ethernet tunnel is established between the host and the access operator NAS. To this end, PPPoE invokes a search protocol that finds the MAC address of the NAS.

With PPPoE, the host initiates a PPP connection. The access operator's NAS terminates the call "momentarily" in order to identify the target ISP for this call. This is done based on the login information the user provides, which is in the form of username@ISP-name. The access operator's NAS then "extends" the PPP connection through a compulsory PPTP/L2TP tunnel ending at the target ISP. Scheme 6 in Fig. 6 shows this case. As in scheme 5, after the user is connected to the ISP, a voluntary PPTP/L2TP tunnel or an IPsec tunnel is established between the user's host and the corporate NAS.

Like L2TP and PPTP, PPPoE is also considered a layer 2 tunneling protocol, because it allows a layer 2 protocol (PPP) to be tunneled over another layer. In fact, at first glance it seems that PPTP/L2TP can replace PPPoE for the tunneling of the PPP connection to the ISP. However, since PPTP/L2TP can be tunneled only over IP, the host must acquire IP connectivity to the access operator NAS before setting up a PPP connection with the ISP. After getting IP configuration from the NAS (e.g., using DHCP), the host can set up a PPP connection over a voluntary PPTP/L2TP tunnel to the selected ISP. The user can now get Internet access through

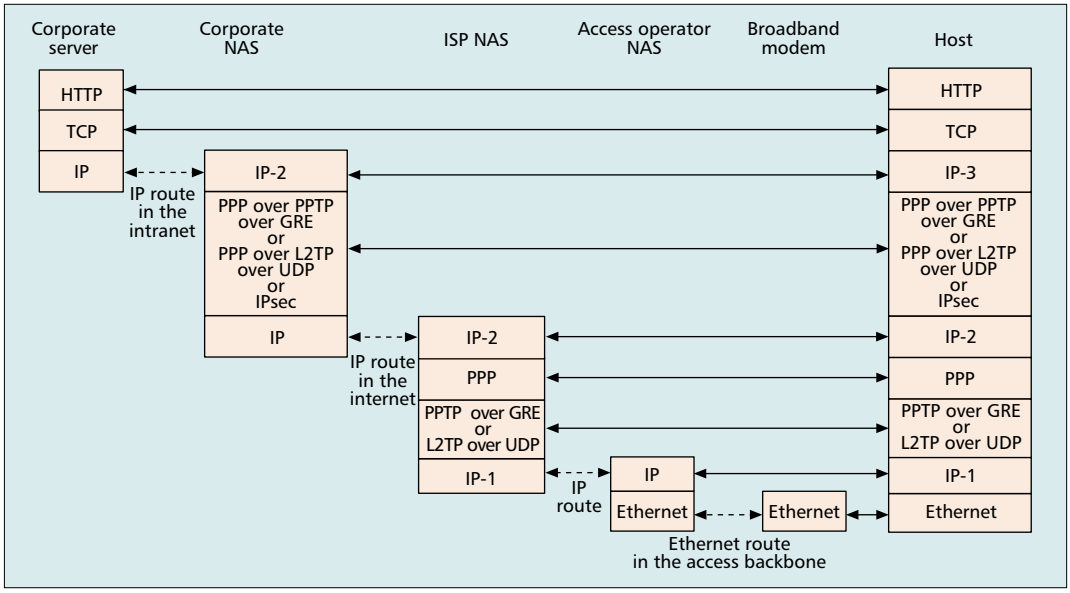

Figure 7. *Using a PPTP/L2TP tunnel rather than a PPPoE tunnel between the host and the ISP's NAS (scheme 7).*

the ISP. Next, a voluntary PPTP/L2TP tunnel or an IPsec tunnel should be set up with the corporate NAS as shown in Fig. 7 (scheme 7). Note that the end host's stack contains in this case three IP layers, and each IP layer is associated with a different pair of source and destination IP addresses as follows:

- IP-1 is the layer that contains the IP address provided by the broadband access operator. Packets carrying this address in their source IP address field can be routed only between the operator NAS and the ISP NAS, but are not permitted to enter the public Internet.
- IP-2 contains the IP address provided by the ISP NAS. Packets carrying this address in their source address field can be routed over the public Internet, but are not allowed to enter the corporate intranet.
- IP-3 contains the source IP address provided by the corporate intranet. Packets carrying this address in their source address field can be routed inside the corporate intranet.

Conclusions

This article discusses the various schemes for establishing an access VPN in a broadband access network, and explains the need for the various tunnels employed in each scheme. We distinguish between two service provisioning approaches in a broadband access networks: the single ISP and multiple ISPs cases. In the first approach the establishment of an access VPN is relatively simple, because the connectivity of the user to its ISP does not require a tunnel, and a tunnel is only needed for VPN connectivity. In the multiple ISPs case we distinguish between two subcases: static service selection, where the association between a user and an ISP is static, and dynamic service selection, where this association can be changed by the user using some online protocol. The first subcase is similar to the single ISP case. However, in the second subcase another tunnel is required between the host and the operator's NAS in order to allow the user to connect to the ISP through PPP. This tunnel can be dispensed with by using an IP-based rather than a PPP-based access, along with a mechanism (e.g., NAT or DHCP FORCERENEW) that allows changing the IP address of the host transparently to the user.

The number of tunnels should be minimized because each tunnel is associated with a processing and bandwidth overhead. In addition, some of the tunnels require the host to get an additional IP address. Our main conclusions are as follows. In the single ISP case and for multiple ISPs with static service selection, a single tunnel between the user and the corporate NAS is sufficient. For multiple ISPs with dynamic service selection, it is recommended to use IP-based access, using NAT or the DHCP FORCERENEW option, in order to be able to establish an access VPN using a single tunnel as well. If the selection of the ISP in the multiple ISPs case is performed using PPP, it is usually better to tunnel the PPP connection using PPPoE rather than L2TP/PPTP, because the former solution requires the host to have only two IP layers in the same stack, whereas the latter solution requires three IP layers. On the other hand, PPPoE requires the operator to use only layer 2 (Ethernet) switches in the access backbone, whereas L2TP/PPTP allows layer 3 (IP) switches as well.

The number of tunnels should be minimized because each tunnel is associated with a processing and bandwidth overhead. In addition, some of the tunnels require the host to get an additional IP address.

References

[1] W. Simpson, "The Point-to-Point Protocol (PPP)," RFC 1331, May 1992.
[2] H. Kummert, "The PPP triple-DES Encryption Protocol (3DESE)," RFC 2420, Sept. 1998.
[3] T. Dierks and C. Allen, "The TLS Protocol Version 1.0," RFC 2246, Jan. 1999.
[4] K. Hamzeh *et al.*, "Point-to-Point Tunneling Protocol (PPTP)," RFC 2637, July 1999.
[5] W. Townsley *et al.*, "Layer Two Tunneling Protocol (L2TP)," RFC 2661, Aug. 1999.
[6] R. Droms, "Dynamic Host Configuration Protocol (DHCP)," RFC 2131, Mar. 1997.
[7] K. Egevang and P. Francis, "The IP Network Address Translator (NAT)," RFC 1631, May 1994.
[8] Y. T. Jones, C. Hublet, and P. De-Schrijver, "DHCP Reconfigure Extension," RFC 3203, Dec. 2001.
[9] L. Mamakos *et al.*, "A Method for Transmitting PPP over Ethernet (PPPOE)," RFC 2516, Feb. 1999.

Biography

Reuven Cohen (M'93, SM'99) received B.Sc., M.Sc., and Ph.D. degrees in computer science from the Technion, Israel Institute of Technology, in 1986, 1988, and 1991, respectively. From 1991 to 1993 he was with IBM T. J. Watson Research Center, working on protocols for high-speed networks. Since 1993 he has been with the Department of Computer Science at the Technion, where he is now an associate processor. He has also consulted for numerous companies, including Hewlett-Packard, ECI Telecom, Terayon, and Innowave, mainly in the context of protocols and architectures for broadband access networks. He serves as an editor of *IEEE/ACM Transactions on Networking* and *ACM/Kluwer Journal on Wireless Networks*.