Ascend

Technical Backgrounder

Multiprotocol Virtual Private Networking Protocols

(L2F, PPTP, ATMP)

ASCEND

Virtual Private Networking Protocols (L2F, PPTP, ATMP)

TABLE OF CONTENTS

OVERVIEW

This document is intended to give the reader an understanding of three existing tunneling protocols. The protocols that will be discussed are Layer Two Forwarding Protocol (L2F), Point to Point Tunneling Protocol (PPTP) and Ascend Tunnel Management Protocol (ATMP).

These protocols are used to provide for Virtual Private Networks (VPN) solutions over the backbone network resources of ISPs and Carriers. They use different methods for tunneling multiple protocol across networks and it is important to understand the differences between protocols when implementing your network. The increased demand for VPNs can be attributed to the shortcomings and high maintenance costs of a privately managed network. As a result, ISPs and Carriers are moving away from the legacy networks that mobile users and remote users access via expensive long distance calls to a more secure and cost-effective (local call) solution using the virtual private networking. They can use virtual private networking to tunnel the different protocols from remote location to the home/central site shared resources. Because there is a shortage of registered IP addresses, there is a demand to use UN-registered IP addresses within the private network and to use a tunneling mechanism to tunnel the traffic across the Internet to the remote locations.

This document is not intended as a competitive analysis, but rather as an comparative guide that describes the three protocols mentioned and delivers a brief comparative analysis. The contents of this document are based on the RFC drafts that have been submitted to the IETF.

Layer Two Forwarding (Protocol) “L2F”

L2F is a protocol that can be used for tunneling link layer protocols such as HDLC, PPP, SLIP, and async HDLC. The LF2F Protocol encapsulates all link layer protocols such that the link layer over the network is totally independent of the user’s link layer protocol. The support of such multiprotocol virtual dial-up application benefits the end user and the ISP because it allows many users to share the large investment in access lines and core network infrastructure and allows users to place local calls use the network. Also, L2F allows for the transport of non-IP protocols in the same secure manner as the rest of the IP infrastructure of the Internet.

The topology consists usually of an ISP Network Access Server (NAS), a Home Gateway and router(s) which guide the tunnel traffic from the NAS to the Home gateway and vice versa.

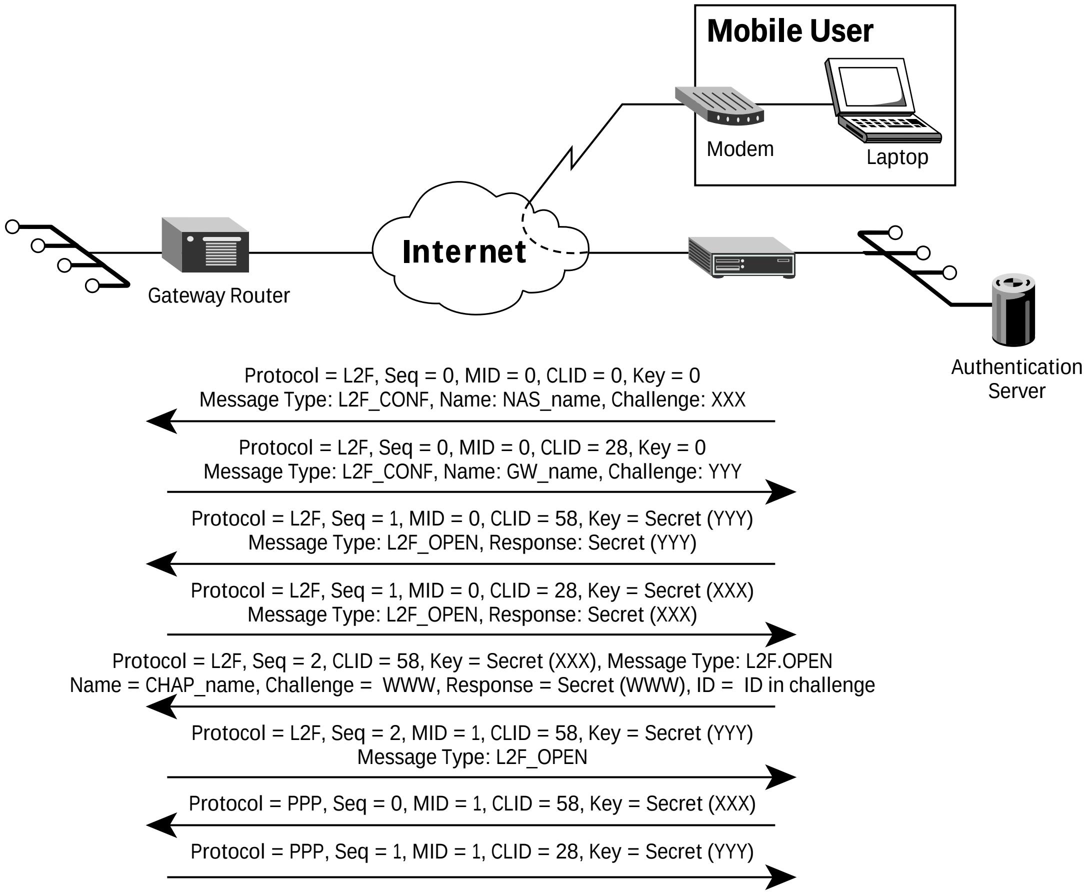
Mobile User
Modem
Laptop
Internet
Gateway Router
Authentication Server
Protocol = L2F, Seq = 0, MID = 0, CLID = 0, Key = 0
Message Type: L2F_CONF, Name: NAS_name, Challenge: XXX
Protocol = L2F, Seq = 0, MID = 0, CLID = 28, Key = 0
Message Type: L2F_CONF, Name: GW_name, Challenge: YYY
Protocol = L2F, Seq = 1, MID = 0, CLID = 58, Key = Secret (YYY)
Message Type: L2F_OPEN, Response: Secret (YYY)
Protocol = L2F, Seq = 1, MID = 0, CLID = 28, Key = Secret (XXX)
Message Type: L2F_OPEN, Response: Secret (XXX)
Protocol = L2F, Seq = 2, CLID = 58, Key = Secret (XXX), Message Type: L2F.OPEN
Name = CHAP_name, Challenge = WWW, Response = Secret (WWW), ID = ID in challenge
Protocol = L2F, Seq = 2, MID = 1, CLID = 58, Key = Secret (YYY)
Message Type: L2F_OPEN
Protocol = PPP, Seq = 0, MID = 1, CLID = 58, Key = Secret (XXX)
Protocol = PPP, Seq = 1, MID = 1, CLID = 28, Key = Secret (YYY)

Figure 1. L2F Tunnel and Client Establishment.

L2F, however, is an unreliable delivery protocol designed to operate over layer two point-to-point links. It is not designed to provide flow control of the data traffic; nor does it provide reliable delivery of the data traffic. Each protocol tunnel carried via L2F is expected to manage the flow control and retry by itself. But the L2F messaging mechanism does have control messages that can be retransmitted. These control messages are expected to be exchanged in lock-step. Thus per-client activities cannot occur until tunnel setup is complete. Therefore, it is very rare to have flow control actions required. If flow control is required, however, sequenced delivery of messages are achieved by using a sequence number in the control messages.

The process begins with the Mobile user initiating a call to the NAS. A point-to-point link is then initiated between the NAS and the Home Gateway, creating the tunnel. The endpoints use Multiplex ID (MID) of zero, prior to allowing any client services. This is used to verify the presence of the remote end, and permit any authentication that needs to take place.

Once the tunnel exists, an unused Multiplex ID (MID) is allocated, and a connect indication is sent to notify the Home Gateway of this new dial-up session. The Home Gateway either accepts or rejects this connection. The initial setup may include authentication information for the Home Gateway to authenticate the user, such as PAP or CHAP, or text dialog for SLIP users.

The authentication of the users takes place in three phases: first at the ISP (NAS); second and an optional third phase at the Home Gateway. In the first phase, the ISP uses the "username" to determine that a virtual dial-up service is required and initiate the tunnel connection to the appropriate Home Gateway. Once a tunnel is established, a new MID is allocated and a session initiated by forwarding the gathered information. For the second phase, the Home Gateway decides whether to accept or reject the connection by validating optional information (PAP, CHAP, etc.) in the dial-up connection. The Optional third phase is with the Home Gateway by pursuing further authentication at the PPP or SLIP layer. This third option is not a requirement of L2F.

Once authentication has completed satisfactorily, the connection is considered to have been "established" by the Home Gateway. A "virtual interface" for SLIP or PPP is created in a manner analogous to what would be used for a direct dialed connection. This virtual interface helps to guide the link layer frames of the tunneled traffic in both directions. Thereafter, frames received from the POP are stripped of any link frame or transparency bytes, encapsulated in the L2F and forwarded over the tunnel. Then the Home Gateway accepts these frames, strips off the L2F and processes them as normal incoming frames for the appropriate interface and protocol. (The "virtual interface" behaves very much like a physical hardware interface.) In the reverse direction, the Home Gateway encapsulates the packet in L2F, and the POP does the stripping of the L2F before transmitting it to the physical interface of the remote user. The packet format for the L2F between the NAS and the Home Gateway is as follow:

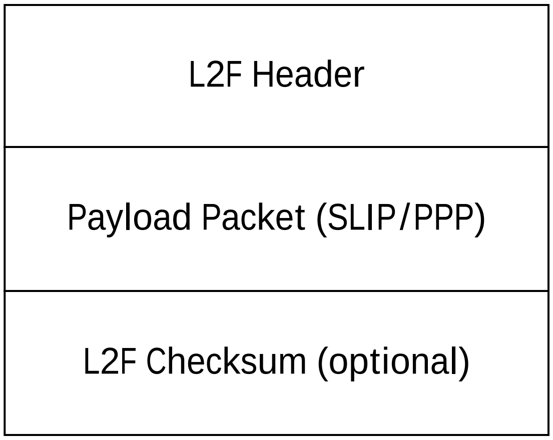

Figure 2. L2F frame format.

At this point the remote user has become just another dial-up client of the home gateway, with the traditional mechanism of authentication and protocol and filtering access. Since the L2F connection notification for PPP clients contain sufficient information for the home gateway to authenticate and initialize the LCP state mechanism, it is not required that the remote user be queried a second time for CHAP authentication. Also, it is not necessary to do multiple LCP negotiations and convergences.

Point to Point Tunnel Protocol “PPTP”

PPTP protocol give users an easy low-cost and secure way to extend a Private Network across the Internet. PPTP tunnels protocol from Front End Processors (FEP) directly into the Windows NT servers and vice versa, encapsulating PPP frame in Enhanced GRE tunnels. (This implementation of PPTP enhanced GRE is backwards compatible with GRE (RFC 1701).) Using PPTP, remote users can employ workstations running Microsoft Windows 95 and Windows NT operating systems to dial into a local Internet service and connect to their corporate network via the Internet backbone transport. To use PPTP, changes to client software are not required; minimal software upgrade is needed for Internet Service Providers. Businesses using PPTP can ensure secure communication by taking advantage of proven authentication and encryption built into Windows NT Remote Access Service.

At first, the Mobile user initiates a call (encrypted in PPP) to the FEP and logs in using their name and password or any other accepted authentication method. The FEP checks with the authentication server (such as radius) to verify the authenticity of the caller. Once the caller has been authenticated, the authentication server returns information necessary for the FEP to build a tunnel across the Carrier/Internet backbone network to the Window NT server.

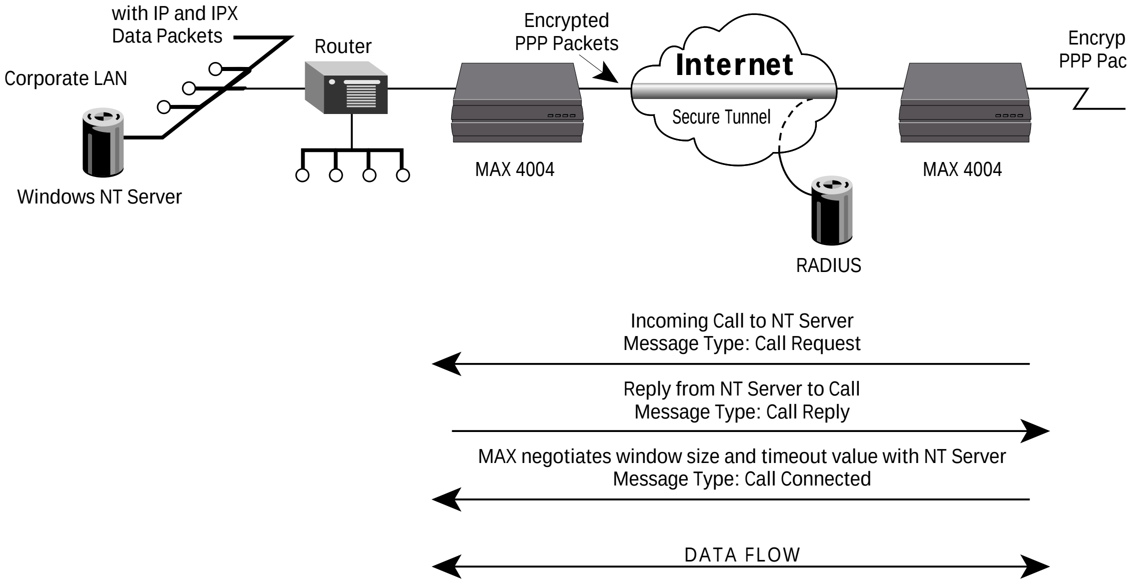

Figure 3. PPTP tunnel data flow.

The FEP then creates an Enhanced Generic Routing Encapsulation (GRE) protocol for use in transporting PPP packets, containing the users data. The format of PPTP tunneled packet is as follow:

Media Header
IP Header
GRE Header |
Payload Packet

Figure 4. PPTP frame structure.

The payload section contains the PPTP payload, which is essentially the PPP packet without the media-specific framing. This means that the frame requires GRE headers that can be transported across IP networks. (This requirement adds 20 bytes of IP plus 16 bytes of GRE, an overhead of 36 more bytes to the header.) The Enhanced GRE allows the packet acknowledgment to be piggybacked on the data packet. However, the acknowledgment can still be sent separately from the data packets.

Once the tunnel is created, PPTP uses its Sliding Window and Adaptive Acknowledgment Time-out features to flow control on each side of the data path, as well as buffering to keep the FEP data channel full. PPTP requires that a time out be used to recover both dropped data and acknowledgment packets. The exact implementation of the time-outs is vendor specific but adaptive time-outs should be implemented with backoff for congestion control. Therefore, there are four key parameters that are used for time-out mechanism and flow control (sliding window) determination:

Packet Processing Delay (PPD) is the delay required for each side to process the maximum amount of data buffered in the receive packet sliding window. This value is exchanged between the FEP and the NTS when a call is connected. For the NTS, this number should be small, but for FEP making modem connection this can be significant.

Sample is the actual amount of time incurred in receiving an acknowledgment for a packet delay already transmitted. This is a measure and not calculated.

Round-Trip Time (RTT) is the estimated round trip for an Acknowledge to be received after a packet has been transmitted. For local network connection this is minimal (or zero), but for Internet this could be large. It would then be necessary to compensate for network transit delay. RTT is adaptive; it will adjust to include the PPD and whatever shifting network delay contributes to the RTT of the packet transitioning across the network.

Adaptive Time-Out (ATO) is the time that must elapse before an acknowledgment is considered lost. After a time-out, the sliding window is partially closed and the timer is backed off.

Therefore, the adjustment of the window is related to the time-outs that occurs. Although each side begins by indicating the size of its received window, it is recommended to use a slow start method to transmit the data initially. In slow start, the transmit window is adjusted until it reaches the maximum window size that was sent by the receive side. As mentioned before, there is no retransmission mechanism. Therefore, once a time-out occurs, the transmission is resumed with the minimum window of one packet, and is adjusted upward with each successful transmission. In case of overflow of the receiver window, all packets are thrown away.

At this point the remote user is like a local client of the Windows NT server. The data packets are flow controlled from the FEP to the Windows NT server as specified above to maintain the integrity of the data.

Ascend Tunnel Management Protocol "ATMP"

ATMP has been developed to provide enhanced virtual private network services. It is based on the TCP/IP protocol suite and uses the Generic Routing Encapsulation (GRE) as described in RFC 1701. ATMP provides a mechanism to dynamically build tunnels through the Internet or across a frame relay network as needed. ATMP is completely transparent to the client dialing into the network as well as the customer's central/home network.

ATMP creates and tears down the tunnel between two Ascend switches or ant other device that supports ATMP. In effect the tunnel collapses the Internet cloud and provides what looks like direct access to a home network. The result of this direct connection is that the packets received through the tunnel must be routed, so ATMP applies only to IP or IPX networks at this time.

A typical network topology of ATMP consist of a mobile node, a foreign agent and a home agent, as well as an authentication server such as Radius.

The Mobile node is a dial in user who wishes to access the corporate shared resources. A mobile dials into a foreign agent.

A Foreign agent is the starting point of the tunnel. The foreign agent requires the radius database to authenticate the mobile user dialing in. The radius data also provides the parameters needed for the foreign agent to start the IP tunnel connection to the home agent.

The Home agent is the terminating point of the remote network. Mobile users often require access to shared resources that possibly utilize unregistered addresses. The home agents accepts/creates the tunnel and authenticates the mobile node.

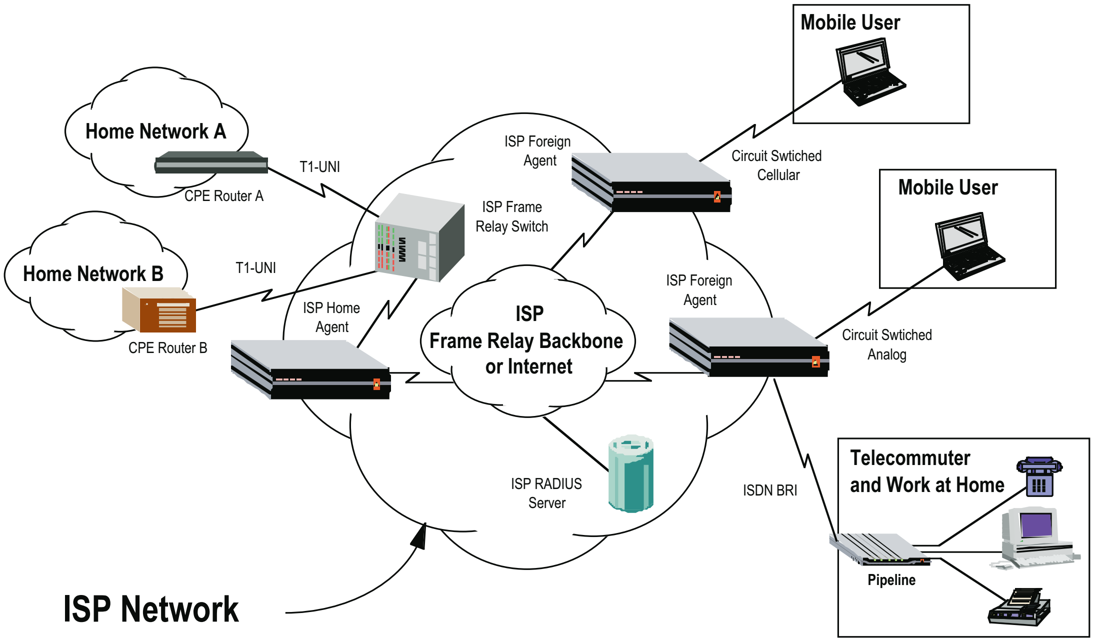

The flow of ATMP connection is as follow :

1. The mobile node dials a connection to the foreign agent.

2. The foreign agent authenticates the mobile node using the radius user profile.

3. The foreign agent locates a connection profile or a radius profile for the home agent based on the attributes in the mobile nodes radius user profile. The profile includes a home agent password that will be used for the mobile user.

4. The foreign agent connects to the home agent using a regular IP connection, which is authenticated in the usual way (such as CHAP).

5. The foreign agent then informs the home agent that the mobile node is connected, and requests that a tunnel be created. The foreign agent can send up to 10 Register request in a two second interval time frame. It will also time-out and log a message if it receives no response to the requests.

6. The home agent requests authentication of the mobile user, by sending a challenge request to the foreign agent. The expected password is the home agent password that was included in the radius profile.

7. The foreign agent sends back a challenge reply to the home agent, which contains the encrypted MD5 password with a given seed.

8. The home agent returns a register reply with a number which identifies the tunnel. If at any point the registration fails, a message is logged and the foreign agent is disconnected. But if the registration is successful a tunnel is created between the foreign agent and the home agent. At this point, the mobile user is attached to the home network as if it was dialed in locally, and data is transferred across the existing tunnel.

9. Once the mobile node disconnects from the foreign agent, the foreign agent sends a deregistration request to the home agent to close the tunnel. This process of deregistration is repeated up to ten times until the foreign agent receives a reply from the home agent. If the foreign agent receives packet for the mobile node whose connection has been closed down, the foreign node discards those packets.

A distinct feature of ATMP is the capability to be able to configure the home agent in either a router or a gateway mode of operation. The difference between the two operation modes is relevant to the relationship of the home agent to the home network.

Router mode: The data from the mobile node is tunneled by the foreign agent to the home agent. The routing module in the home agent forwards the packets onto the local network. The network may be the home network, or may it may support another router that can connect to the home network. Regardless, the packet delivery does rely on the routing mechanism such as RIP or static routes to reach the final destination and not the WAN connections. In the case of routing an IPX packet from the mobile node, the home agent must see the mobile node as connected to another IPX network. This virtual IPX network is added to the home agents routing table automatically based on the IPX attributes it receives from the foreign agent. The mobile node user profile must specify the IPX network number that is unique with in the enterprise.

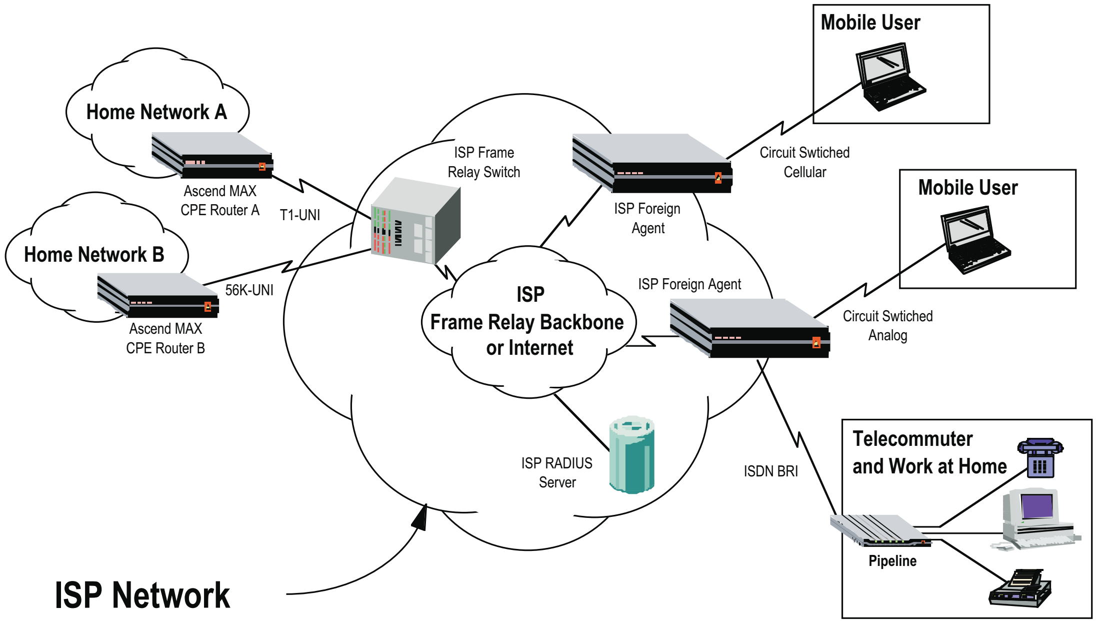

Gateway mode: The data from the mobile node is tunneled by the foreign agent to the home agent. Then the packets that are received by the home agent are tunneled through to the home network across an open WAN connection. The WAN connection must be on-line. The home agent will not bring up a WAN connection to the home network based on a packet received through the tunnel. For this reason, ATMP gateway mode is usually used when the connection is nailed (leased line) or configured in such a way so the WAN is not brought down by the idle timer.

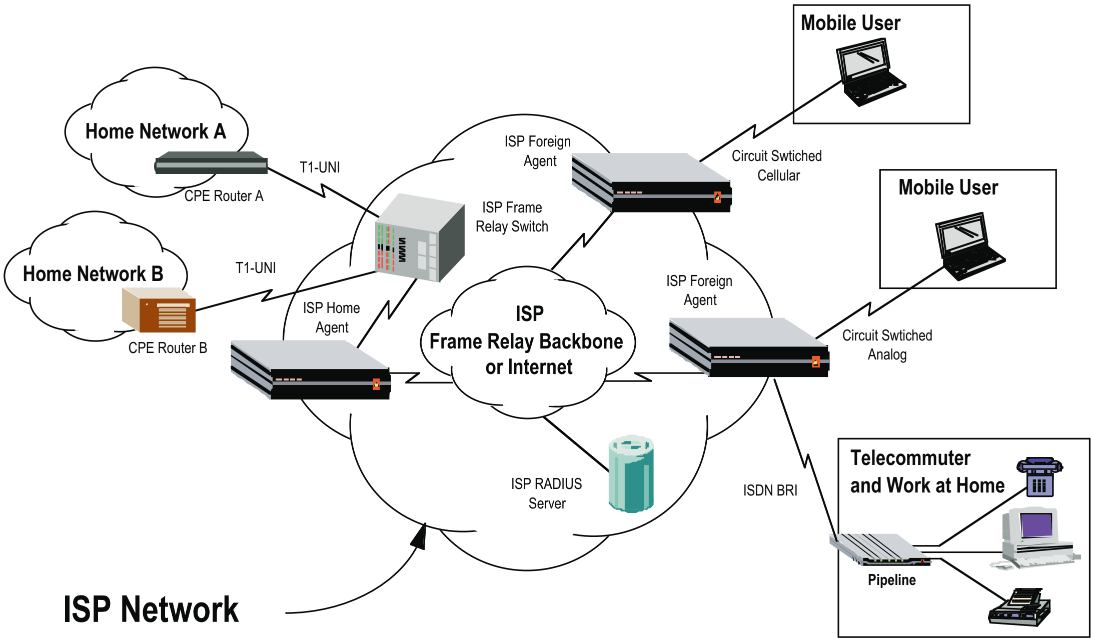

Regardless of the Home agents mode of operation, for each tunnel between the home agent and the foreign agent, there is a distinct tunnel identifier. The UDP port that the home agent listens to for receiving messages, is set to default to port number 5150. This port assignment is configurable. The UDP port assignment for units to listen to are also configurable and have the range of 0 to 65535.

From a security point of view the registration function of the ATMP is protected by the challenge/respond mechanism which is similar to CHAP. The home agent challenges each registration mechanism attempted by the foreign agent. Therefore the authentication requires configuration of a shared secret for each Home Agent/Foreign Agent pair.

In either Gateway/Router mode of operation of the home agent, from the mobile users point of view, it is as if the Mobile user is locally connected to the shared resources at the home network. From the network managers view point, all the data traffic is securely tunneled from the local foreign agent to the home agent and routed or forwarded across a WAN connection to the final destination securely.

Comparative summary
ATMP vs. PPTP :

ATMP and PPTP are both GRE-based protocols. However, ATMP is the UDP/IP based routing protocol for dynamic management of GRE tunnels. ATMP tunnels are created between two agents (foreign and home) with no Windows/NT requirements. The Home agent can be configured in two different modes of operation (router or gateway) which allows network managers to control the tunneled traffic received from the foreign agent in a manner suited for the home network. The home network can either be local (router mode), or remote via a nailed WAN link (Gateway mode).

PPTP, on the other hand, requires that the mobile users have Window NT or Window 95, or be PPP clients. In addition, the home agent must be a Windows NT server, where the tunnel terminates in PPP frames. The PPTP tunnel utilizes an enhanced GRE tunnel which controls the flow of traffic between the two endpoints of the tunnel. PPTP also supports NetBios while ATMP tunnels IP and IPX. Additionally, PPTP encrypts PPP packets prior to tunneling across the network. ATMP only encrypts the authentication passwords using MD5 while data is passed in the clear.

ATMP vs. L2F :

The L2F protocol is similar in some ways to ATMP. The key differences is the encapsulation method used for creating and managing the packets across the tunnel. L2F is not based on GRE (RFC 1701) but rather a new proposed proprietary mechanism based on PPP. To support L2F, the mobile client must run PPP. The home gateway and the Network Access Server (local site) must run L2F. The advantage of L2F is that it supports transporting protocols other than IP and IPX such as Appletalk, which ATMP does not provide. However, L2F requires each home network to have its own dedicated home agent/gateway, which can be expensive.

On the other hand, ATMP is more dynamic in managing the GRE tunnels between the home agent and the foreign agent. The ATMP tunnels are created between two agents (foreign and home) and do not require a dedicated home agent for each segment. The home agent can be configured in two different modes of operation (router or gateway) which allows management of the received tunneled traffic in a manner suitable to the home network.

PPTP vs. L2F

To compare PPTP & L2F we need to look at :

- Technology for transport across the tunnel
- Interaction with the application of the Operating System
- Network protocol communication through the tunnel

The protocol L2F utilizes a different encapsulation method for creating and managing the packets across the tunnel. L2F is not based on GRE (RFC 1701) but rather a new proposed mechanism which is based on PPP. To support L2F, the mobile client must run PPP, and the Home gateway and the Network Access Server (local site) must run L2F. In contrast with PPTP, L2F provides support for Appletalk and does not require a Windows NT server to front end each Home Network.

On the other hand, L2F is not really designed as an end-to-end technology. Although PPTP requires a Window NT server to front end the Home Network, the GRE tunnel terminates in PPP frames. The PPTP tunnel is an enhanced GRE tunnel with flow control mechanisms for both endpoints of the tunnel. As a result, the PPTP tunnel never gets saturated and is very simple to implement since it is based on TCP. In addition, the PPTP payload header is much smaller since it does not require or use the tunnel Identifier.

In contrast, L2F does not provide a flow controlling mechanism for the data across the tunnel; rather, it relies on each of the protocols that are encapsulated to manage the flow control and retransmission that is necessary. The only flow control mechanism supported by L2F is for management of the tunnel and the L2F management frames. Therefore, L2F does not make any assumption about the underlying media, whereas PPTP assumes IP fabric for the creation of the GRE tunnel.

Both tunnel protocols impose different constraints on the network protocols communicating through the tunnels. L2F passes the CHAP negotiated username, challenge and the credentials at the time of indicating the connection. In contrast, PPTP does not specify user identity in the protocol, since passing credential requires the FEP to be trusted. Otherwise, any FEP on the network can pick-up those credentials and bombard the server until one is found that the server accepts.

Also, L2F does not have support for all the different types of credentials, such as MS, CHAP, KAP, EAP; PPTP does have the support for all types of credentials. In most secure networks, one or more of those credentials commonly are used, and there are other installed based equipment that utilize them. Therefore, having the support for the various types of credentials is of great importance.

Another security aspect that PPTP has taken into account is that it allows only LCP negotiations to be done end-to-end, since LCP states cannot be assumed by the servers. In contrast, L2F allows servers to pick up LCP states even though the servers were not participants in the LCP negotiation phase. While L2F does not have support for Callback and ISDN/Modem pooling scenarios, PPTP does.

ASCEND

05-13

Ascend. Remote Networking Solutions That Work.™

Ascend Communications, Inc.
1275 Harbor Bay Parkway
Alameda, CA 94502, USA
TEL: 510.769.6001
FAX: 510.814.2300

E-mail: info@ascend.com
Toll Free: 800.621.9578
FAX Server: 415.688.4343
Internet Home Page:
http://www.ascend.com