Project Report
ECE 646 (Fall 2001)

Comparison of VPN Protocols – IPSec, PPTP, and L2TP

Poonam Arora, Prem R. Vemuganti, Praveen Allani

Department of Electrical and Computer Engineering
George Mason University
Fairfax, VA 22202

Table of Contents

List of Figures

1 Introduction

1.1 What is Virtual Private Network (VPN)?

First of all, it is a ***network,*** that is, it provides inter-connectivity to exchange information among various entities that belong to the VPN. Secondly it is ***private***, that is it has all the characteristics of a private network. So, "what characterizes a private network?" A private network supports a closed community of authorized users, allowing them to access various network-related services & resources. The traffic originating & terminating within a private network traverses only those nodes that belong to the private network. Further, there is traffic isolation. It means that, the traffic corresponding to this private network does not affect nor is it affected by other traffic extraneous to the private network. The final characteristic of a VPN is that it is ***virtual.*** A virtual topology is built upon an existing, shared physical network infrastructure. A Virtual Private Network (VPN) is the extension of a private network that encompasses links across shared or public networks like the Internet. A VPN enables you to send data between two computers across shared or public internetworks in a manner that emulates the properties of a point-to-point private link. The act of configuring & creating a virtual private network is known as virtual private networking[6].

To emulate a point-to-point link, data is encapsulated, or wrapped, with a header that provides the routing information allowing it to traverse the shared or public internetworks to reach its endpoint. To emulate a private link, the data being sent is encrypted for confidentiality. Packets that are intercepted on the shared or public network are indecipherable without the encryption keys. The link in which the private data is encapsulated or encrypted is known as a virtual private network (VPN) connection.

A VPN [3] connection allows users working at home or on the road to connect in a secure fashion to a remote corporate server using the routing infrastructure provided by public internetworks (such as the Internet). From the user's perspective, the VPN connection is a point-to-point connection between the user's computer & a corporate server. The nature of the intermediate internetworks is irrelevant to the user because it appears as if the data is being sent over a dedicated private link.

VPN connection also allows a corporation to connect to branch offices or to other companies over a public internetwork (such as the Internet), while maintaining secure communications. The VPN connection across the Internet logically operates as a wide area network (WAN) link between the sites. In both these cases, the secure connection across the internetwork appears to the user as a private network communication—despite the fact that this communication occurs over a public internetwork. Hence, the name - *virtual private network.*

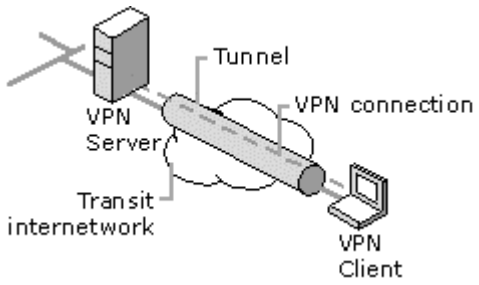

Figure 1Virtual Private Network Connection

1.2 Why VPNs are so popular today?

Traditional private networks facilitate connectivity among various network entities through a set of links, comprising of dedicated circuits (T1, T3 etc). These are leased from public telecommunication carriers like MCI-WorldCom or Regional Bell Operating Companies (RBOCs) as well as privately installed wiring. The capacity of these links is available at all times, albeit fixed & inflexible. The traffic on these private networks belongs only to the enterprise or company deploying the network. Therefore, there is an assured level of performance associated with the network. Such assurances come with a price. These can be viewed as[6]:

- Traditional private networks are not cheap to plan & deploy. The costs associated with dedicated links are especially high when they involve international locations. The planning phase of such networks involves detailed estimates of the applications, their traffic patterns and their growth rates. Also, the planning periods are long because of the work involved in calculating these estimates.
- Further, dedicated links take time to install. It is not unusual that telecommunication carriers take about 60 to 90 days to install & activate a dedicated link Such a long waiting period adversely affects the company's ability to react to the quick changes in the these areas.
- Another recent trend is the mobility of today's work force. Portable computing facilities such as laptops & palm-based devices have made it easy for people to work without being physically present in their offices. This also makes less investment into real estate.
- To support the increase in home offices, companies need to provide a reliable IT infrastructure so employees can access company information from remote locations. This has resulted in large modem pools for employees to dial-in remotely. The cost keeps increasing due to the complexity of managing & maintaining the large modem pools. An additional cost with the mobile users is the long-distance calls or toll-free numbers paid for by the company. The costs are much higher if we consider international calling. For, companies with large, mobile workforce, these expenses add up to significant numbers. Also, dial-in connections limit the remote user to a maximum access speed of 56Kbps for analog modems & 128Kbps for Integrated Services Digital Network (ISDN).

These limitations hamper the day-to-day activities that require high-speed access to the Intranet as available from a regular office. Even though a high-speed access media like cable modems & Digital Subscriber Lines (DSL) can overcome the access limitations, the service providers offering this cannot have easy access to a company's Intranet due to Firewall & Security restrictions.

So, there is an urgent need for a reliable mechanism to authenticate valid users & restrict their accesses based on their access privileges. Hence a Virtual Private Network (VPN) can help resolve many of the above issues associated with today's private networks. The advantages of a VPN are:

- A VPN facilitates an agile IT infrastructure.
- Global VPNs enable connectivity to all locations anywhere in the world at a fraction of the cost of dedicated links.
- VPN services enable remote access to the Intranet at significantly lower cost, thus enabling support for a mobile workforce.
- Additionally, the VPN architecture supports a reliable authentication mechanism to provide easy access to the Intranet from anywhere using any available access media including analog modems, ISDN, cable modems, ISDN, DSL & wireless.

1.3 Types of VPN Services

There are three basic types of VPN services:
1. Local Area Network (LAN) Interconnect VPN services.
2. Dial-up VPN services.
3. Extranet VPN services.

1.3.1 LAN Interconnect VPN

LAN Interconnect VPN services help to interconnect local area networks located in multiple geographic areas over the shared network infrastructure. Typically, this service is used to connect multiple geographic locations of a single company. Several small offices can be connected with their regional & main offices. This service provides a replacement for the expensive dedicated links. The advantage is the flexibility, like:

- It's easy to increase the capacity of any of the links depending upon the applications supported on the VPN.
- As applications change with time, the architecture can be adapted to meet the demands.
- Additional geographical sites can be connected to the VPN with little effort[6].

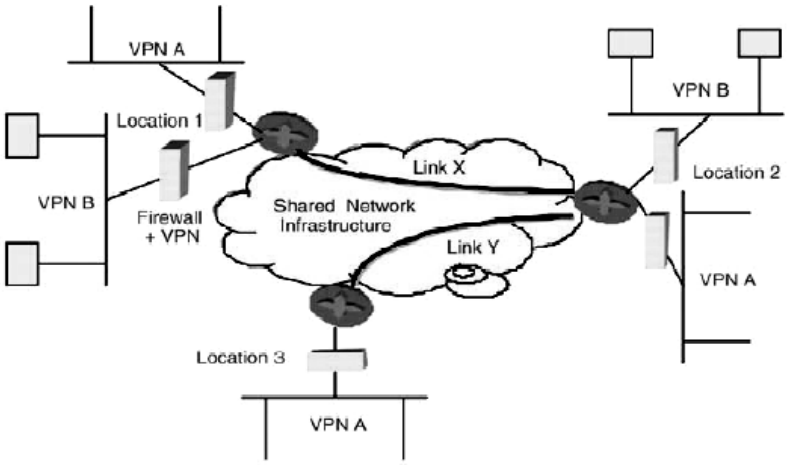

Figure 2 LAN Interconnect

In the above figure both VPNs A & B is implemented on top of a shared network infrastructure. Another advantage is reduction in cost. Dedicated private links are expensive. Also, using a shared infrastructure is cheap because of the economies of scale. The costs are borne by the different VPNs that are supported on the infrastructure. In Figure 2, the links X and Y are borne by VPNs A and B, which in turn reduces the costs of Companies A and B. But, still the characteristics of the VPN are retained even though a shared network infrastructure is used, by providing mechanisms to isolate & secure the traffic of each VPN.

1.3.2 Dial-up VPN Services

It helps mobile & telecommuting employees in accessing the company's Intranet from remote locations. The remote employee (user) dials into the nearest Remote Access Server (RAS), which

establishes a secure connection to the company's Intranet, usually through a firewall enhanced with VPN capabilities. Upon successful authentication of the user, the secure connection enables the user to have access to the company's propriety.

The above is the case with one dial-up VPN model, called the Layer 2 Tunneling Protocol (L2TP), usually aimed at providing services to the home offices & telecommuters who dial-in to a specific local RAS.

Alternatively, Point-to-Point Tunneling Protocol (PPTP) model focuses on the mobile user, who may dial-in to any local ISP. The user initiates a connection to any of the VPN servers located in the company's Intranet, after establishment of the connection with the ISP. Then, the access privileges are established after the authentication server validates the user. In this model RAS does not participate in establishment of the VPN connection. So, RAS configuration is not need with PPTP model. PPTP can also be used with high-speed accesses like DSL & cable modems.

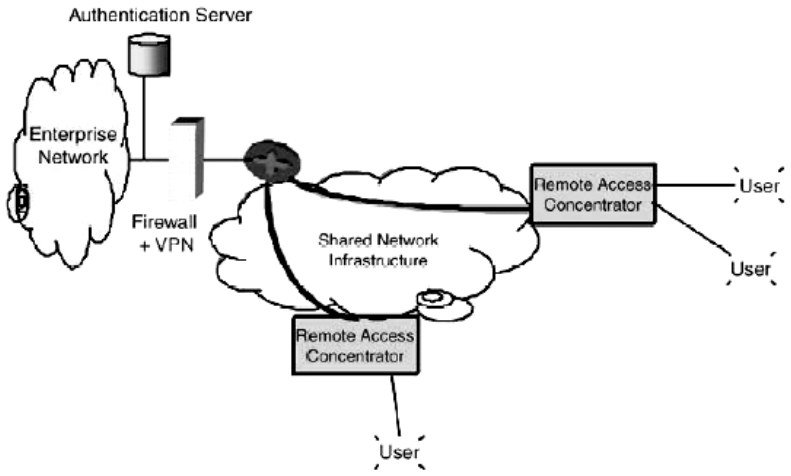

Figure 3 Dial-Up VPN Services

The advantages of the dial-up VPN services are:

- It results in considerable cost-savings to a company & eliminates the managing of large modem pools & uses the RAS that belong to local ISP's.
- The dial-in VPN service takes advantage of high-speed access, thus, eliminating some access limitations of the home-office.

1.3.3 Extranet VPN Services

An Extranet VPN service combines LAN interconnect & dial-in VPN services. Figure 4 shows the typical VPN.

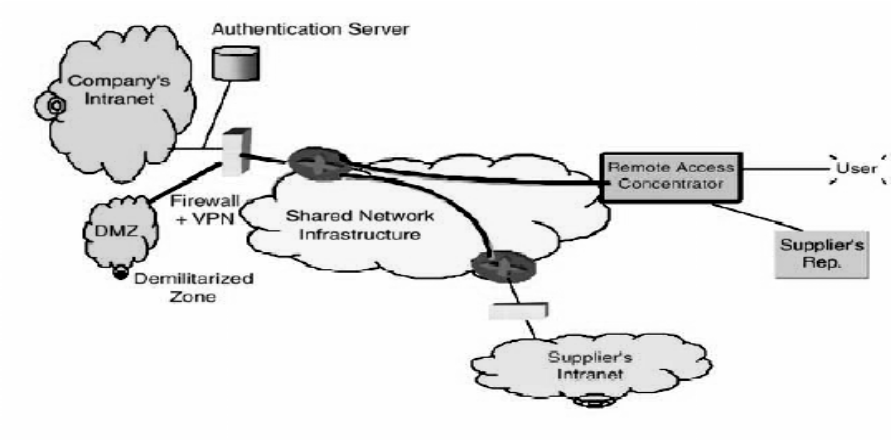

Figure 4 Extranet VPN Service

This infrastructure enables external vendors, suppliers & customers to access specific areas of the company's Intranet. The specific area is denoted as Demilitarized Zone (DMZ). When a suppliers representative connects to the company's Intranet or dialing in remotely, the firewall & authentication mechanism ensure that the connection is directed to DMZ. A Company's employee on the other hand has full access to the company's Intranet.

The advantages of extranet services are:

- The flexibility of it helps provide connectivity to new external suppliers & customers within a short period of time.
- The fast communications supported by the extranet helps in several e-commerce areas including efficient inventory management & electronic data interchange (EDI), which helps in significant savings, in cost & effectively compete in the rapidly growing market.

2 Overview of Tunneling

The process of using an internetwork infrastructure to transfer data for one network over another network is called Tunneling. The data to be transferred can be frames (or packets) of another protocol. The tunneling protocol encapsulates the frame in an additional header, instead of sending frame as it is produced by an originating node. The additional header is required for providing the routing information to the encapsulated payload to traverse the intermediate internetwork [8].

The encapsulated packets are then routed between tunnel endpoints over the internetwork. The logical path through which the encapsulated packets travel through the internetwork is called a "tunnel". Once these encapsulated packets reach the destination they are decapsulated to get the original data. Tunneling includes this entire process encapsulation, transmission, and decapsulation of packets.

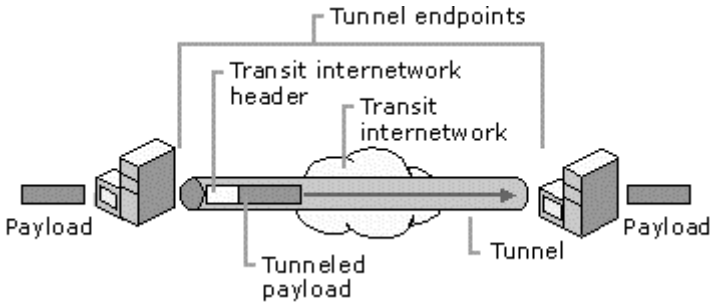

Figure 5 Tunneling Technique

The transit internetwork can be any internetwork. The most widely known real world example is the public internetworks "the Internet". While Internet provides one of the most pervasive and cost-effective internetworks, any other private or public internetwork that acts as a transit internetwork can replace it.

2.1 Tunneling Technologies

Some of the mature tunneling technologies include:

- SNA tunneling over IP internetwork
- IPX tunneling for Novell NetWare over IP internetwork
- Point-to-Point tunneling protocol
- Layer 2 tunneling protocol
- IPSec tunnel mode

For a tunnel to be established, both the tunnel client and tunnel server must be using the same tunneling protocol. The tunneling functionality can be based on:

- Layer 2 tunneling protocol: This corresponds to the data-link layer & use frames as their unit of exchange. PPTP & L2TP operate on this layer & both encapsulate the payload in a PPP frame to be sent across an internetwork.
- Layer 3 tunneling protocol: This corresponds to the Network layer & use packets. IPSec tunnel mode operates on this layer & it encapsulates the packets in an additional header before sending them across an IP internetwork.

And, these layers correspond to the Open Systems Interconnection Reference Model (OSI).

2.2 How Tunneling works for different protocols?

For Layer 2 Tunneling technologies, such as PPTP & L2TP

- A tunnel is similar to a session
- Both tunnel endpoints must agree to the tunnel & must negotiate configuration variables.
- Data transferred across the tunnel is sent using a datagram-based protocol & tunnel maintenance protocol is used as the means to manage the tunnel.
- Also, a tunnel must be created, maintained & then terminated.

For Layer 3 tunneling technologies, such as IPSec

- It's assumed that all the configuration issues are preconfigured often by manual process.
- No tunnel maintenance is required.

Once the tunnel is established, tunnel data can be sent. The tunnel client or server uses a tunnel data transfer protocol to ready the data to be transferred.

For example, when the tunnel client sends a payload to the tunnel server or vice versa:

- The tunnel data transfer protocol header is added to the payload
- The client then sends the resulting encapsulated payload across the internetwork, which routes it to the tunnel server.
- The tunnel server accepts the packets, removes the tunnel data transfer protocol header, and forwards the payload to the target network.

2.3 Tunneling Protocols and Requirements

Layer 2 tunneling protocols; PPTP & L2TP are based on well-defined PPP protocol with useful features. These useful features & their Layer 3 counterparts; IPSec, address the basic VPN requirements as given below:

- **User Authentication:** Layer 2 tunneling protocols inherits the user authentication schemes of PPP, including the EAP methods. Many Layer 3 tunneling schemes generally assume that the endpoints are well known and authenticated before the tunnel is established. But, IPSec uses Internet Key exchange (IKE) negotiation, which provides mutual authentication of the tunnel endpoints is an exception.
- **Token Card Support:** Using the Extensible Authentication Protocol (EAP), layer 2 tunneling protocol can support a wide variety of authentication methods, including one-time passwords, cryptographic calculators, and smart cards. IPSec uses similar methods like, public key certificate authentication in its IKE negotiation.
- **Data encryption:** Layer 2 tunneling protocols support PPP-based data encryption mechanisms. For encryption these protocols use generally RSA/RC4 algorithm. Layer 3 tunneling protocols can use similar methods like, IPSec. It uses generally, DES, 3DES,Blowfish, RC5 algorithms for encryption.
- **Key Management:** Layer 2-encryption mechanism usually relies on initial key generated during user authentication & refreshes it periodically. IPSec explicitly negotiates a common key during the IKE exchange, and also refreshes it periodically.
- **Multiprotocol support:** Layer 2 tunneling supports multiple payload protocols, which makes it easy for tunneling clients to access their corporate networks using IP, IPX, NetBEUI, and so on. In contrast, Layer3 tunneling protocols, such as IPSec tunnel mode, typically support only target networks that use the IP protocol.

3 IP Security Architecture

IPSec is designed to provide interoperable, high quality, cryptographically based security for Internet protocol version 4 (IPv4) and Internet Protocol version 6 (IPv6). The set of security services offered includes access control, connectionless integrity, data origin authentication, protection against replays (a form of partial sequence integrity), confidentiality (encryption), and limited traffic flow confidentiality. These services are provided at the IP layer, offering protection for IP and/ or upper layer protocols[1][5][7].

IPSec provides security services at the Internet Protocol (IP) layer by enabling a system to select required security protocols, determine the algorithm(s) to use for the service(s) and put in place any cryptographic keys required to provide requested services. IPSec can be used to protect one or more "paths" between a pair of hosts, between a pair of security gateways, or between a security gateway and a host. The term "security gateway" refers to an intermediate system (router or a firewall) that implements the IPSec protocols.

The security services are provided through the use of two security protocols-the Authentication Header (AH) and the Encapsulating Security Payload (ESP)- and through the use of cryptographic key management procedures and protocols, including the Internet Security Association and Key Management Protocol (ISAKMP) and the Internet Key Exchange protocol (IKE).

The AH header comes right after the IP header and contains cryptographic hashes of data and identification. The AH protects the source and destination addresses of the IP header.

The ESP header allows for encryption of the data payload protecting data privacy and integrity. ESP works with symmetric encryption algorithms: DES, 3DES and Blowfish. In order to use IPSec you need to have the same protocols, encryption algorithms and keys on both sides of the connection.

IKE provides the mechanism for two IPSec entities to negotiate security services and their associated session authentication and encryption keys.

3.1 Security Associations

The concept of a security association (SA) is fundamental to the IP security architecture. An SA defines the kinds of security measures that should be applied to the packets based on who is sending the packets, where they are going, and what type of payload they are carrying. The set of security services offered by a SA depends on the security protocol and its options and the mode in which the SA operates. SAs can be negotiated dynamically between two communicating peers when they wish to use one or more of IPSec's security services, based on the security policies given by the security administrator. Alternatively, albeit more rarely, SAs can be statically specified by the administrators directly.

A SA is uniquely identified by three parameters: a destination IP address, a security protocol identifier, and a Security Parameter Index (SPI). The destination IP address is the IP address of the destination endpoint for the SA. The SPI is a 32-bit number usually chosen by the destination endpoint of the SA, and it has local significance only within that destination endpoint. The security protocol identifier is the protocol number for either AH (51) or ESP (50).

Note that the source IP address is not used to define a SA. This is because a SA is a security services agreement between two hosts or gateways for data sent in one direction. As a result, if two peers need to exchange information in both directions using IPSec, two SAs are required, one for each direction.

SAs operate in two modes: transport mode and tunnel mode. Transport mode is designed primarily to protect the higher-layer protocols (e.g., TCP and UDP). In tunnel mode, an IP packet becomes the payload for another IP packet. This allows the inner IP packet, including its IP header, to be subjected to encryption or other security measures, whereas the outside IP packet serves to steer the data through the network. Hosts can provide both transport mode and tunnel modes, whereas security gateways can provide only tunnel mode (unless the gateway is acting as a host, in which case it can provide either mode).

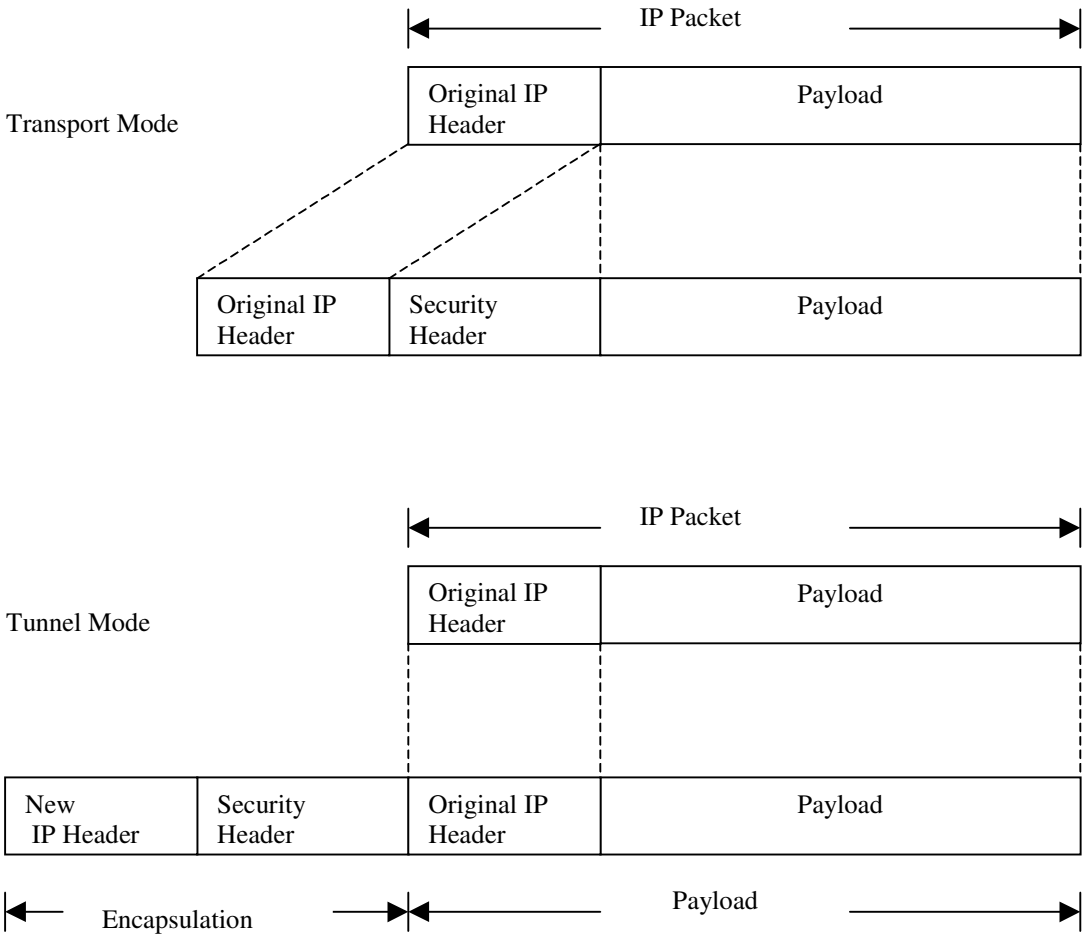

Figure 6 Transport and Tunnel Modes

Figure 6 shows the IP packet structures for transport and tunnel modes. In transport mode, the original IP header remains mostly intact, and a security header is placed between the IP header and its payload. The original IP header is modified only to the extent that it now reflects that a security header follows rather than the payload. In tunnel mode, the original IP packet becomes

the payload of an encapsulating IP packet. The encapsulating IP header indicates that a security header follows.

A SA can be viewed as a unidirectional channel, offering either AH or ESP security. Note that each SA can only offer one of these security services. If both AH and ESP protection is applied to a data stream, then two SAs must be established and maintained. Similarly, to secure bi-directional communications between two hosts or security gateways, two SAs (one in each direction) are required. The term SA bundle is applied to a sequence of SAs through which traffic must be processed to satisfy a specific security policy.

3.2 Security Databases

Two databases are associated within an IPSec node: the Security Policy Database (SPD) and the Security Association Database (SAD). A policy administrator composes a set of security policies to meet the security needs of all types of IP traffic both into and out of this node. These policies are kept in the SPDs to be used in directing the processing of IP packets and the construction of SAs as needed. All SAs are registered in the SAD, along with the SAs' parameters.

3.2.1 Security Policy Database

The SPD is constructed with the policies that specify what services are to be offered, i.e., what addresses have IPSec applied at what standard of security, and what addresses are passed through without IPSec.

The protection offered is based on requirements defined by a security policy database (SPD) established and maintained by a user or system administrator (or by an application operating within constraints established by either of the above). In general, IP packets are selected for one of the three processing modes based on IP and transport layer header information matched against entries in the SPD. Each packet is either afforded IPSec security services, discarded, or allowed to bypass IPSec security services entirely.

3.2.2 Security Association Database

The SAD contains parameters associated with each security association (SA) that has been determined with the SPD.

3.3 Authentication Header

IPSec Authentication Header protocol (AH) is used to provide per-packet authentication- that is, data integrity and data origin authentication for the IP payload (upper-layer protocol header and data) and as much of the IP header as possible. Depending on which cryptographic algorithm is used and how keying is performed, the AH may also provide non-repudiation of origin services. Finally, the AH may offer an anti-replay service at the discretion of the receiver, to help counter specific denial-of-service attacks.

The fundamental mechanism used by AH to provide authentication is the authentication header. The new IP packet is formed by placing the authentication header either after a new IP header or a slightly modified original IP header. The authentication header contains six fields as shown in figure 7.

Next Header (8 bits)	Payload Length (8 bits)	Reserved (16 bits)
Security Parameters Index (SPI) (32 bits)		
Sequence Number (32 bits)		
Authentication Data (n * 32 bit)		

32 bit

Figure 7 The Authentication Header (AH) format

Next Header The Next Header field is an 8-bit field that identifies the type of protocol header immediately after the AH.

Payload length This 8-bit field specifies the length of AH in 32-bit words, minus 2.

Reserved This 16-bit field is reserved for future use, and is set to 0.

Security Parameters Index This 32-bit field together with the IP destination address, identifies the Security association. The SPI is an arbitrary value chosen upon establishing the SA. The value 0 may be used to indicate that a SA has not yet been established, while the values 1-255 are reserved for future use.

Sequence Number This 32-bit field indicates the packet's position in sequence. The sequence number is initialized to 0 upon establishing a SA, and is incremented by one with each sent packet. If replay protection is enabled, this value must never be allowed to cycle. When all sequence number values have been used, a new SA and thus a new key must be established.

Authentication Data This field contains the Integrity Check Value (ICV) generated by using an authentication algorithm and a cryptographic key specified in the corresponding SA. The length of this field is not fixed, and depends on the algorithm specification. The sender computes the data prior to sending the IP packet, and the receiver verifies it upon reception.

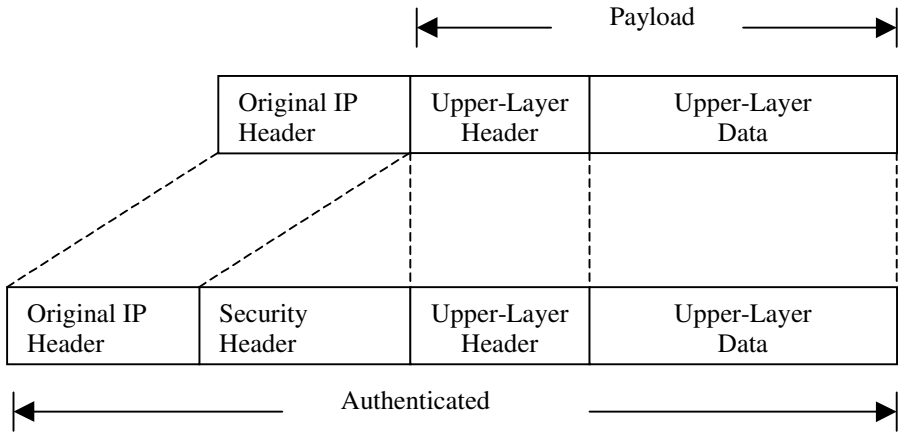

Figure 8 AH Transport mode

Transport Mode

In transport mode, the original IP header is retained as the header of the new IP packet, and the authentication header is inserted between the IP header and the original payload, as shown in Figure 8.

Transport mode has the advantage of adding only a few extra bytes to the original IP packet. However, because the original IP header is used as the header of the new IP packet, only end hosts can use AH transport mode. This limitation is acceptable when both endpoints of the IPSec SA are acting on behalf of other devices. Another disadvantage is that transport mode can be subjected to the traffic pattern analysis, in which an observer gleans information from the number and types of packets traversing the network even if the contents are obscured.

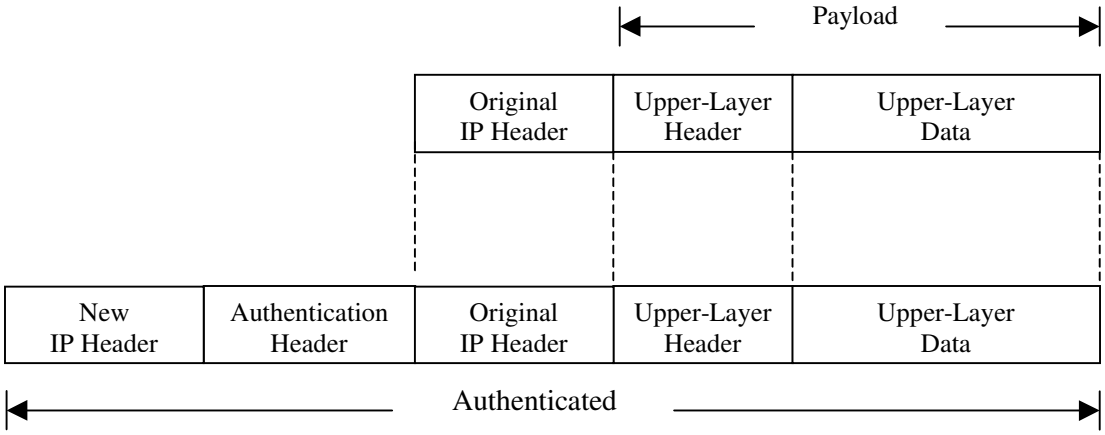

Figure 9 AH Tunnel mode

Tunnel Mode

In tunnel mode, a new IP header is created for the new IP packet, and the authentication header is inserted between the original and new IP headers, as shown in Figure 9. The original IP packet stays intact and is encapsulated within the new IP packet. In this way, authentication is provided over the entire original IP packet, in addition to the authentication header and the immutable fields of the new IP header. Further, more processing power is needed for adding and stripping these extra headers.

The original IP header is completely unaltered and contains the ultimate destination as well as the original source IP address. The new IP header contains the source and destination IP addresses of the IPSec devices between which the new packet will travel. Consequently, tunnel mode can be used whether the endpoints of the SA are hosts or security gateways.
If the SA is between hosts, the new source and destination IP addresses are usually the same as the original. The typical reason for using AH in tunnel mode between hosts is to completely authenticate the original packet.

If the SA is between security gateways, the new source and destination IP addresses are those of the gateways. Tunnel mode AH between security gateways allows aggregation of traffic between sites through an authenticated tunnel. In addition, traffic analysis is more difficult because the

original IP packet is buried within the payload of the new IP packet and potentially is multiplexed with other traffic traveling between the same sites.

AH is algorithm independent, which means that AH will operate with the algorithm of choice, depending on the level of security required. Currently the algorithm options are HMAC (hashed message authentication code), MD5 (message digest 5) or HMAC SHA1 (Secure Hash Algorithm).

Optionally AH will, if selected, provide protection against replays (man-in-the-middle attacks) as long as the receiver checks the sequence numbers. AH authenticates all of the packets including the upper protocol data, with the exception of the destination address. AH can be used alone, when only authentication is required, or in combination with ESP when a higher level of security is required.

3.4 Encapsulating Security Payload

The IPSec Encapsulating Security Payload protocol (ESP) provides authentication, data confidentiality through encryption, and optional anti-replay protection for IP packets.

As with AH, additional fields are inserted into an IP packet to provide these services, and many of the fields have the same meaning as in AH. Unlike with AH, however, these fields are spread throughout the IP packet. Some are in an ESP header, others are in an ESP trailer, and one is in an ESP authentication segment as shown in Figure 10.

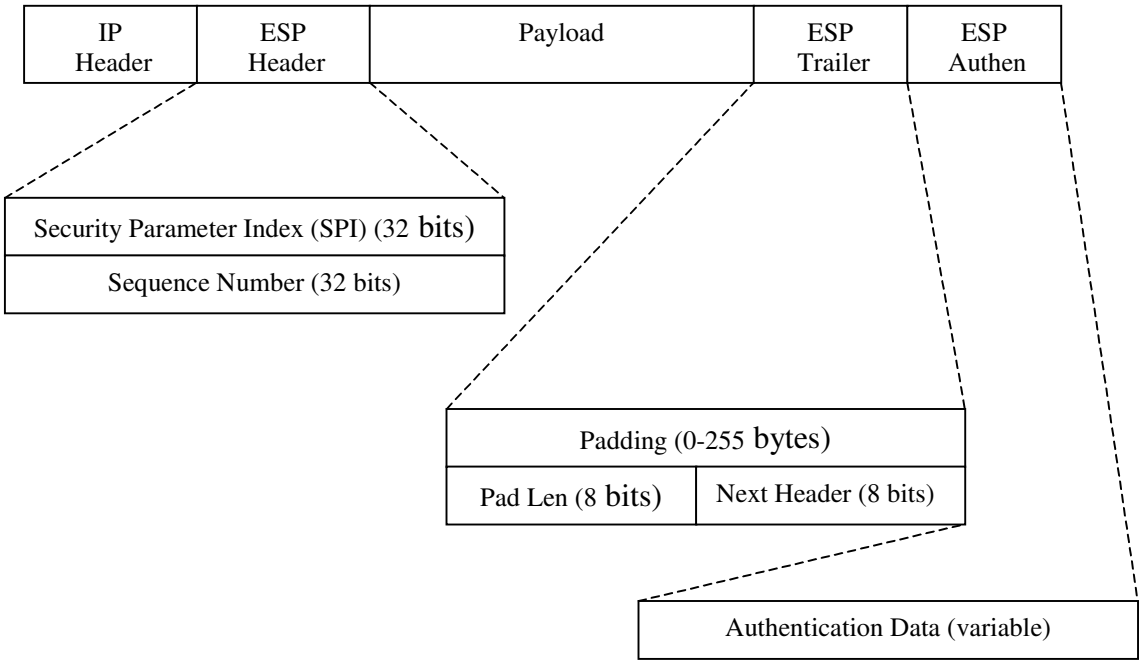

Figure 10 ESP header, trailer, and authentication segment formats

The ESP header follows either a new IP packet or a slightly modified original IP header, depending on the mode. The ESP trailer follows the end of the original IP packet, and the ESP authentication segment follows the trailer. If authentication is not applied, the ESP authentication segment is not appended. If encryption is applied, everything from the end of the ESP header to the end of the ESP trailer is encrypted.

The format of the Encapsulating Security Payload is shown in Figure 11. The fields within the ESP header, trailer, and authentication segments are similar with those within the authentication header. In fact, the SPI, Sequence Number, Next Header, and Authentication Data fields are defined just as they are in the AH protocol. The additional fields are described below.

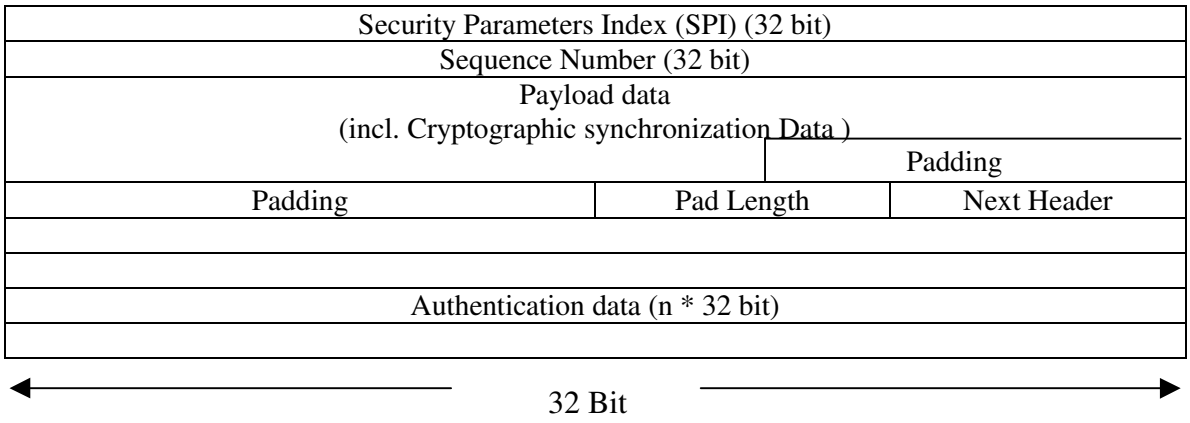

Figure 11 The Encapsulating Security Payload (ESP) format

Payload Data This field is the actual data carried by the packet, and the type of data is indicated by the Next Header field. Any cryptographic synchronization data required by the encryption algorithm, e.g., an Initialization Vector (IV), may be carried explicitly in the Payload field. An encryption algorithm that requires such explicit, per-packet synchronization data must indicate the length, any structure for such data, and the location of this data.

Padding The padding field contains 0-255 randomly generated types of data, and can be used for several purposes:

- Some encryption algorithms require the cleartext to be a multiple of some number of bytes, in which case the Padding field is used to fill out the cleartext (consisting of Payload Data Pad Length and Next Header Fields) to the size required by the algorithm.
- Padding may also be required, irrespective of encryption algorithm requirements, to ensure that the resulting encrypted payload terminates on a 4-byte/32-bit boundary as required by the ESP packet format.
- Padding may also be used to conceal the actual length of the payload, in support of (partial) traffic flow confidentiality.

Pad Length The Pad length field indicates the number of pad bytes immediately preceding it. The range of valid values 0-255, where a value of zero indicates that no Padding bytes are present. The Pad Length field is mandatory.

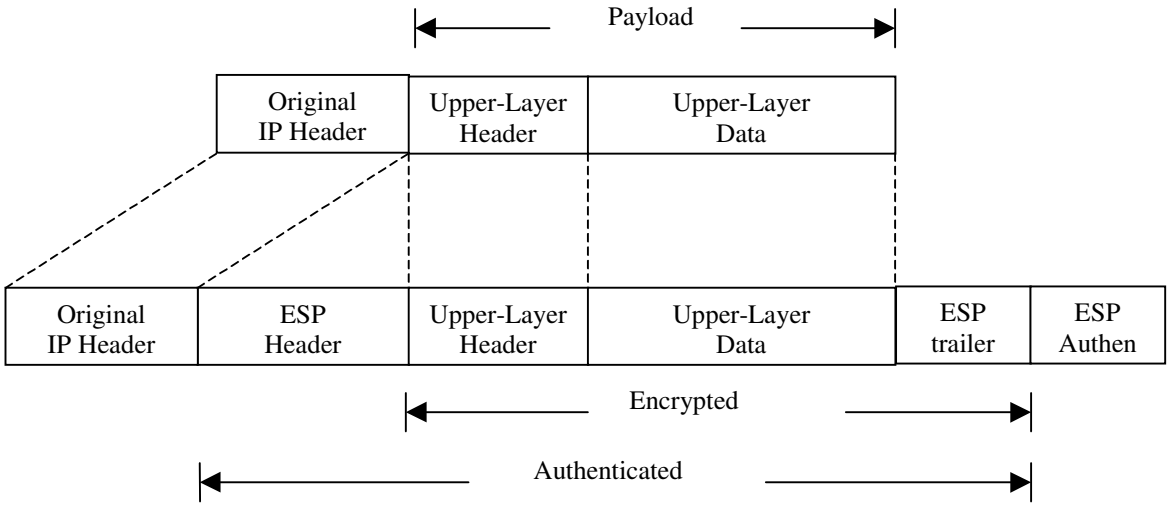

Figure 12 ESP Transport mode

Transport Mode

Figure 12 shows how the new IP packet is constructed from the original IP packet be inserting the ESP header between the IP header and payload, and by appending the ESP trailer and authentication segments, if necessary. If the original IP packet already had IPSec security headers, the new ESP header is placed before them. Because the original IP header is used, the IP packet's source and destination address cannot change. Therefore, ESP in transport mode, like AH in transport mode, can be used only between hosts.

Transport mode is most useful when it is not necessary to hide or authenticate the source and destination IP addresses.

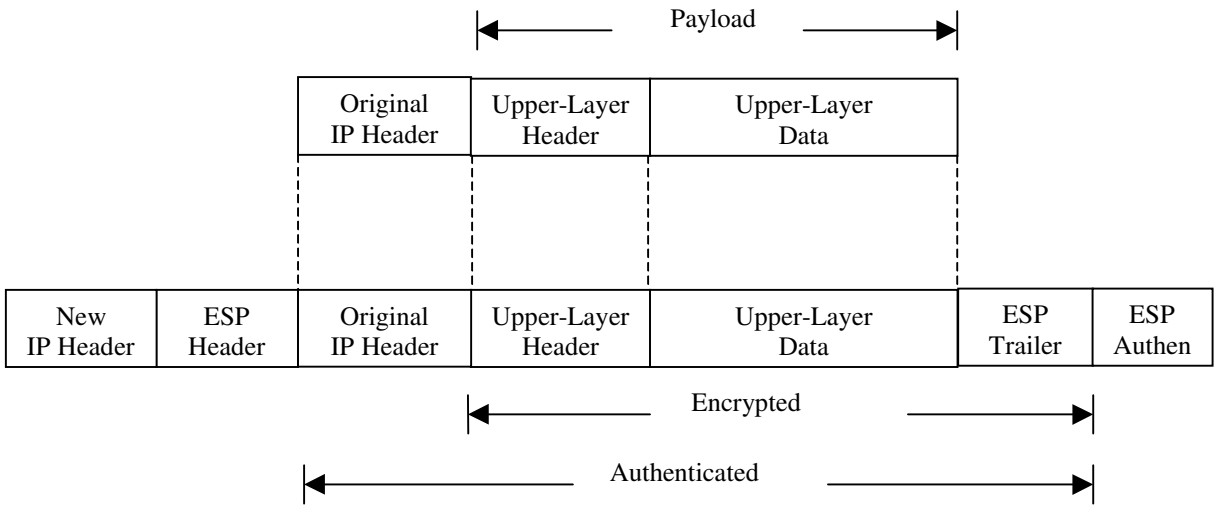

Figure 13 ESP Tunnel mode

Tunnel Mode

Tunnel mode encapsulates the entire original IP packet into the new IP packet. Figure 13 shows how a new IP header and an ESP header are added to the beginning of the original IP packet, and the ESP trailer and authentication segments are appended to the end. If the tunnel is between hosts, the source and destination IP addresses in the new IP header can be the same as the addresses in the original. If the tunnel is between two security gateways, the addresses in the new IP header reflect the gateway addresses. Running ESP in tunnel mode between security gateways can provide both confidentiality and authentication for the transit traffic between the two gateways.

ESP is also designed to be algorithm independent and the options are DES, 3DES, RC5, Blowfish, Idea, Cast, and others are being added.

	Data Origin Authentication	Data Integrity	Replay Protection	Data Confidentiality
AH	Yes	Yes	Yes	No
ESP	Yes	Yes	Yes	Yes

Table 1: Comparison of security services provided by AH and ESP Headers

3.5 Internet Key Exchange

Although the IPSec ESP and AH protocols specify how the data security services are to be applied to each IP packet according to the SAs negotiated between the IPSec devices, they do not tell how these SAs are actually negotiated. The SAs can be manually configured by the system administrators or, more importantly, they can be dynamically negotiated via a key management protocol such as Internet Key Exchange.

IKE is based on the framework defined by the Internet Security Association and Key Management Protocol (ISAKMP). It implements part of the Oakley and SKEME (Secure Key Exchange Mechanism) key exchange methods, as well as two exchange methods of its own.

3.5.1 ISAKMP

The **ISAKMP** defines the procedures for authenticating a communicating peer, creation and management of Security Associations, key generation techniques, and threat mitigation (e.g., denial of service and replay attacks). All of these are necessary to establish and maintain secure communications (via IP Security Service or any other security protocol) in an Internet environment. It is not linked to one specific algorithm, key generation method, or security protocol.

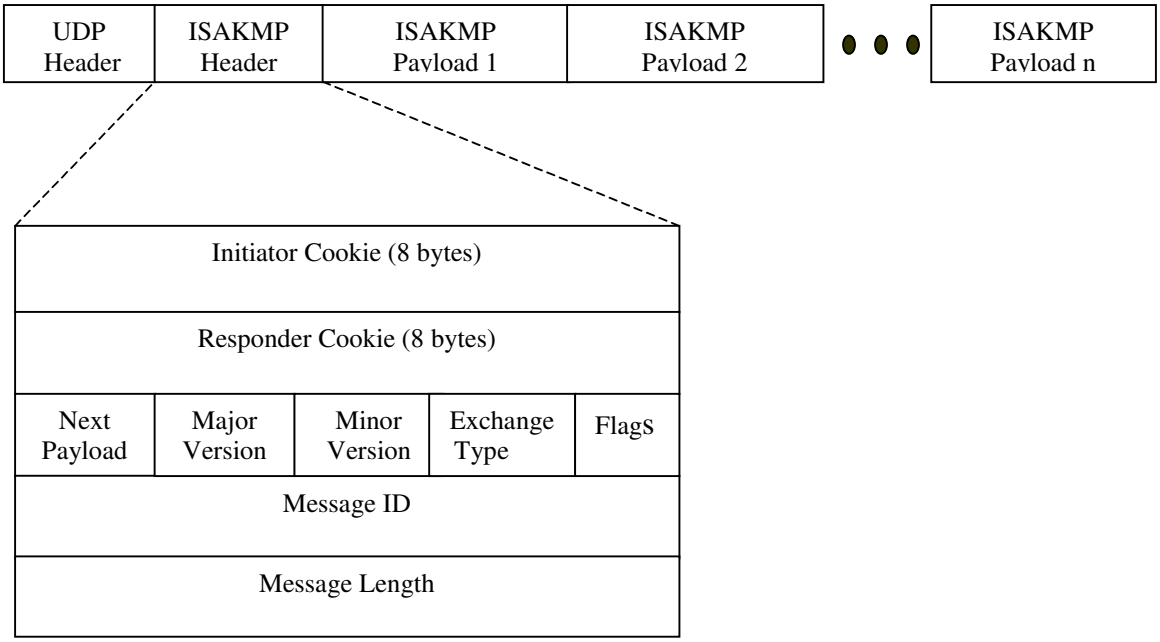

Figure 14 ISAKMP message format

An ISAKMP message consists of an ISAKMP header and one or more ISAKMP payloads chained together in a UDP (port 500) packet. This is shown is Figure 14. The initiator cookie and the responder cookie are special values generated by the ISAKMP peers to provide some protection against denial-of-service attacks, in which an attacker might try to generate lot of bogus ISAKMP messages and overwhelm the ISAKMP processor. A cookie affords some protection by permitting ISAKMP to discard bogus messages quickly before too many of the processing resources have been wasted. The two cookies are also used for identifying the security association between the two ISAKMP peers after the negotiation has been completed successfully.

ISAKMP defines two phases in the security association negotiation. The first phase is the negotiation between the two ISAKMP peers. In this phase, two peers agree on how to protect further communications between them, thus establishing an ISAKMP security association. An ISAKMP security association should not be confused with the IPSec SA. An ISAKMP SA is bi-directional, and does not actually apply to the IPSec traffic.
In the second phase, security associations for other protocols (IPSec in particular, although in theory an ISAKMP SA can be used by other protocols) are negotiated between the two ISAKMP peers. Because the communication channel between the negotiators is already secure, these subsequent negotiations can proceed more quickly. In many cases,

3.5.2 Oakley

Oakley is a key determination protocol, which uses Diffie-Hellman key exchange algorithm. Oakley supports Perfect Forward Secrecy (PFS), which ensures that if a single key is compromised, it permits access only to data protected by a single key. It never reuses the key that protects communications to compute additional keys and never uses the original key-generation material to compute another key.

Diffie-Hellman (DH) Technique

The Diffie-Hellman Technique (named for its inventors Whitfield Diffie and Martin Hellman) is a public key cryptography algorithm that allows two communicating entities to agree on a shared key. Diffie-Hellman starts with the two entities exchanging public information. Each entity then combines the other's public information along with its own secret information to generate a shared-secret value.

3.5.3 SKEME

SKEME constitutes a compact protocol that supports a variety of realistic scenarios and security models over Internet. It provides clear tradeoffs between security and performance as required by the different scenarios without incurring in unnecessary system complexity. The protocol supports key exchange based on public key, key distribution centers, or manual installation, and provides for fast and secure key refreshment. In addition, SKEME selectively provides perfect forward secrecy, allows for replaceability and negotiation of the underlying cryptographic primitives, and addresses privacy issues as anonymity and repudiatability.

IKE protocols present different exchanges in modes, which operate in one of two phases. Phase 1 establishes an ISAKMP Security Association and derives shared secrets, which are used to protect phase 2 exchanges. Phase 2 negotiates security associations on behalf of IPSec or other security services, and generates fresh keying material. There are three basic modes: main mode and aggressive mode, used in phase 1 and quick mode, used in phase 2.

Main mode uses an exchange of six different messages between two IPSec endpoints to complete negotiation of authentication of the endpoints and keying material. This negotiation, if required, will provide perfect forward secrecy (PFS), which means that, after the first two messages are exchanged, subsequent communication is protected.

Aggressive mode authenticates the endpoints with only three messages, but it does not provide PFS. Moreover, SA negotiation is limited with aggressive mode.

Quick mode is used after the tunnel is established to regenerate fresh key material for encryption purposes- it does not authenticate the endpoints. The new key data is used to encrypt subsequent communications data, which is why we indicated earlier that 56 bit DES is often strong enough.

4 PPTP

A consortium, consisting of Ascend Communications, 3Com, ECI Telematics, U.S. Robotics, and Microsoft, developed the PPTP specification for the tunneling of data across the Internet.

The PPTP protocol [4] is built on the well-established Internet Communications Protocol PPP (point-to-point protocol), and TCP/IP (Transmission Control Protocol/Internet Protocol). Multiprotocol PPP offers authentication as well as methods of privacy and compression of data. IP is routable, and has an Internet infrastructure. PPTP allows a PPP session to be tunneled through an existing IP connection, no matter how it was set up. An existing connection can be treated as if it were a telephone line, so a private network can run over a public one.

Tunneling is achieved because PPTP provides encapsulation by wrapping packets of information (IP, IPX, or NetBEUI) within IP packets for transmission through the Internet. Upon receipt, the external IP packets are stripped away, exposing the original packets for delivery. Encapsulation allows the transport of packets that will not otherwise conform to Internet addressing standards.

PPTP encapsulates Point-To-Point Protocol (PPP) frames into IP data grams for transmission over an IP-based Internet work, such as Internet. To encapsulate PPP frames as tunneled data, PPTP uses a TCP connection known as PPTP control connection to create, maintain and terminate the tunnel & a modified version of Generic Routing Encapsulation (GRE).

PPTP inherits encryption or compression or both, of PPP payloads from PPP. Authentication that occurs during the creation of PPTP-based VPN connection uses the same authentication mechanisms as PPP connections, such as:

- Extensible Authentication Protocol (EAP)
- Challenge Handshake Protocol (CHAP)
- Shiva Password Authentication Protocol (SPAP), and
- Password Authentication Protocol (PAP)

PAP

Password Authentication Protocol (PAP) provides a method for the peer to establish its identity using a 2-way handshake. This is one of the ways of user authentication. A stronger authentication such, as CHAP must negotiate prior to PAP.

GRE

The Protocol GRE (Generic Routing Encapsulation) is for performing encapsulation of an arbitrary network layer protocol over another arbitrary network layer protocol.

The payload is first encapsulated in a GRE packet, which possibly includes the route. The resulting GRE packet can then be encapsulated in some other protocol and then forwarded. The outer protocol is the delivery protocol.

4.1 Tunneling in PPTP

A tunnel is defined by PNS (PPTP network server) - PAC (PPTP access concentrator) pair. The tunnel protocol is defined by a modified version of GRE .The tunnel carries PPP datagrams between the PAC and the PNS. A control connection operating over TCP controls the establishment, release and maintenance of sessions and of the tunnel itself.

Before PPP tunneling can occur between a PAC and PNS, a control connection must be established between them. The control connection is a standard TCP session over which PPTP call control and management information is passed.
This tunnel is used to carry all user session PPP packets for sessions involving a given PNS-PAC pair.

4.2 Types of Tunneling:

Tunnels can be created in various ways:

4.2.1 Compulsory Tunneling

Compulsory tunneling (also referred to as NAS-initiated tunneling) enables users to dial to NAS (Network Access Servers), which then establishes tunnel to the server. The connection between the client of the user and the NAS is not encrypted.

4.2.2 Voluntary Tunneling

Voluntary tunneling (also referred to as client-initiated tunneling) enables clients to configure and establish encrypted tunnels to tunnel servers without an intermediate NAS participating in the tunnel negotiation and the establishment.

For PPTP, only voluntary tunneling is supported.

PPTP works by encapsulating the virtual private network packets inside of PPP packets which are in turn encapsulated in Generic Routing Encapsulation (GRE), packets sent over IP from the client to the gateway PPTP server and back again. In conjunction with this encapsulated data channel, there is a TCP-based control session. The control session packets are used to query status and convey signaling information between client and the server. The control channel is initiated by the client to the server on TCP port. In most cases this is a bi-directional communication channel where the client can send requests to the server and vice-versa.

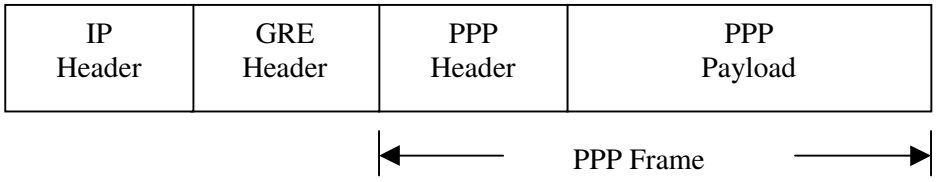

Figure 15 Structure of a PPTP packet containing user data

Tunnel Maintenance with the PPTP Control Connection

The PPTP control connection is between the IP address of the PPTP client (using a dynamically allocated TCP port) and the IP address of the PPTP server (using the reserved TCP port 1723). The PPTP control connection carries the PPTP call control and management messages that are used to maintain the PPTP tunnel. This includes the

transmission of periodic PPTP Echo-Request and PPTP Echo-Reply messages to detect a connectivity failure between the PPTP client and PPTP server. PPTP control connection packets consist of an IP header, a TCP header, and a PPTP control message as illustrated in Figure 16.

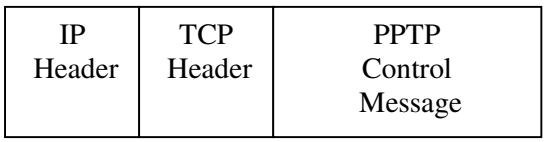

Figure 16 PPTP Control Connection Packet

PPTP Data Tunneling

PPTP data tunneling is performed through multiple levels of encapsulation.
Figure 17 shows the resulting structure of PPTP tunneled data.

The initial PPP payload is encrypted and encapsulated with a PPP header to create a PPP frame. The PPP frame is then encapsulated with a modified GRE header. GRE was designed to provide a simple, lightweight, general-purpose mechanism for encapsulating data sent over IP internetworks. GRE is a client protocol of IP using IP protocol 47.

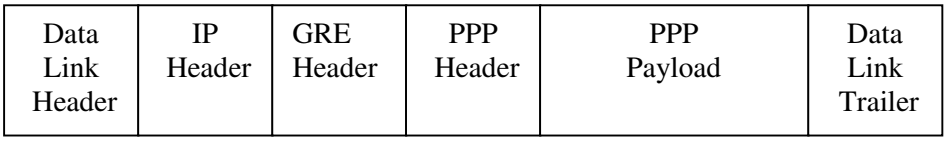

Figure 17 PPTP Tunneled data

4.3 Microsoft PPTP

Microsoft has implemented its own algorithms and protocols to support PPTP. This implementation of PPTP, called Microsoft PPTP, is used extensively in commercial VPN products precisely because it is already a part of the Microsoft Windows 95, 98, and NT operating systems.

The authentication protocols in Microsoft PPTP are the Microsoft Challenge/Reply Handshake Protocol version 1 (MS-CHAPv1) and version 2 (MS-CHAPv2); the encryption protocol is Microsoft Point to Point Encryption (MPPE). MS-CHAPv2 is available as an upgrade for Microsoft Windows 95, Windows 98, and Windows NT 4.0. Even though this upgrade is available, we believe that most implementations of PPTP use MS-CHAPv1.
MS-CHAP version 2 is highly recommended as they provide mutual authentication and are the most secure methods of exchanging credentials.

4.3.1 Authentication in PPTP

The Microsoft PPTP server can only be run under Windows NT, although client software exists for windows NT, Windows some, and Windows 98. There are three authentication options supported in the Microsoft implementations:

1) Clear Password: The client sends the server a password in the clear.
2) Hashed Password: The client sends the server a hash of the password.
3) Challenge/Response: The client and the server authenticate using the MS-CHAP challenge/response protocol.

Microsoft Windows NT uses two one-way hash functions to protect passwords: the LAN Manager hash and the Windows NT hash. The LAN Manager hash is based on the DES encryption algorithm the Windows NT hash is based on MD4 one-way hash function. Both of these hash functions are used in PPTP.

MS-CHAP version 1

MS-CHAP version 1 works as follows:

1. Client requests a login challenge.
2. Server sends back an eight-byte random challenge.
3. The client calculates the LAN Manager hash, and adds five nulls to create a 21-byte string, and partitions the string into three seven-byte keys. Each key is used to encrypt value. This is returned to the Server as a response. The client does the same with Windows NT hash.
4. Server looks up the hash in its database, encrypts the challenge with the hash, and compares it with the encrypted hashes it received. If they match, the authentication completes.
5. The server could make comparison on the Window NT hash or the LAN Manager hash; the results would be the same. Which hash the server uses depends on a particular flag in the packet. If the flag bit is set, the server tests against the Windows NT hash; if the flag bit is not set, the server tests against the LAN Manager hash.

4.3.2 Encryption in PPTP

Microsoft Point-to-Point Encryption (MPPE) may be used with PPTP to provide an encrypted connection but PPTP itself doesn't use encryption. MPPE uses the RC4 algorithm with either 40- or 128-bit keys. All keys are derived from the cleartext authentication password of the user. RC4 is stream cipher; therefore, the sizes of the encrypted and decrypted frames are the same size as the original frame.

RC4 is stream cipher designed by Ron Rivest for RSA data security. It is a variable key-size stream cipher with byte-oriented operations. The algorithm is the use of random permutation. Analysis shows that the period of the cipher is overwhelmingly likely to be greater than 10^{100} years. Eight to sixteen machine operations are required per o/p byte; and the cipher can be expected to run very quickly in software. Independent analysts have scrutinized the algorithm and it is considered secure.

For VPNs, IP data grams sent across the Internet can arrive in a different order from the one in which they were sent, and a higher proportion of packets can be lost. Therefore, MPPE for VPN connections changes the encryption key for each packet. The decryption of each packet is independent of the previous packet. MPPE includes a sequence number in the MPPE header. If packets are lost or arrive out of order, the encryption keys are changed relative to the sequence

number. Although this level of encryption is satisfactory for many applications, it is generally regarded as less secure than some of the encryption algorithms.

There are two modes of encryption in MPPE:

- Stateful
- Stateless

Stateful encryption will provide the best performance but may be adversely affected by networks experiencing substantial packet loss. If you choose stateful encryption you should also configure flow control to minimize the detrimental effects of this lossiness.

Because of the way that the RC4 tables are reinitialized during stateful synchronization, it is possible that two packets may be encrypted using the same key. For this reason, Stateful encryption may not be appropriate for lossy network environments (such as Layer 2 tunnels on the Internet)

Stateless MPPE Encryption provides a lower level of performance, but will be more reliable in a lossy network environment.

5 Layer 2 Tunneling Protocol (L2TP)

Layer Two Tunneling Protocol (L2TP) [2] is a combination of Microsoft's PPTP & Layer 2 Forwarding (L2F), a technology proposed by Cisco System's, Inc. L2TP supports any routed protocol such as, IP, IPX, and AppleTalk. It also supports any WAN technology including frame relay, ATM, X.25, and SONET. L2TP can be used as a tunneling protocol over the Internet or private Intranets.

PPP defines an encapsulation mechanism for transporting multiprotocol packets across layer 2 (L2), point-to-point links. Typically, a user obtains a L2 connection to a Network Access Server (NAS) using one of a number of techniques (dial-up, ISDN etc) and then runs PPP over that connection. In such a configuration, the L2 termination point and PPP session endpoint reside on the same physical device (i.e., the NAS)

L2TP extends the PPP model by allowing the L2 and PPP endpoints to reside on different devices Interconnected by a packet-switched network. With L2TP, a user has an L2 connection to an access-concentrator and the access-concentrator then tunnels individual PPP frames to the NAS. This allows the actual processing of PPP packets to be divorced from the termination of the L2 circuit.

L2TP uses UDP messages over IP internetworks for both tunnel maintenance and tunneled data. L2TP therefore uses message sequencing to ensure the delivery of messages. L2TP supports multiple calls for each tunnel. To identify the tunnel and a call, there is a Tunnel ID and Call ID in the L2TP control message and the L2TP header for tunneled data.

Authentication that occurs during the creation of L2TP tunnels must use the same authentication mechanisms as PPP connections such as, EAP, CHAP, SPAP, and PAP.

5.1 Types of Tunneling

L2TP is used in two different scenarios:

1. Compulsory Tunneling
2. Voluntary Tunneling

5.1.1 Compulsory Tunneling

It refers to the scenario in which a network node- a dial or network access server, for instance acting as LAC, extends a PPP session across a backbone using L2TP to remote LNS. This. This implies that the NAS must be preconfigured to know the tunnel's terminating endpoint given a particular user's authentication information. This is done without the explicit intervention of the remote user, and the remote computer needs no special software- this operation is transparent to the user initiating the PPP session to the LAC. This allows for the decoupling of the location and/or ownership of the modem pools used to terminate dial calls, from the site to which users are provided access.

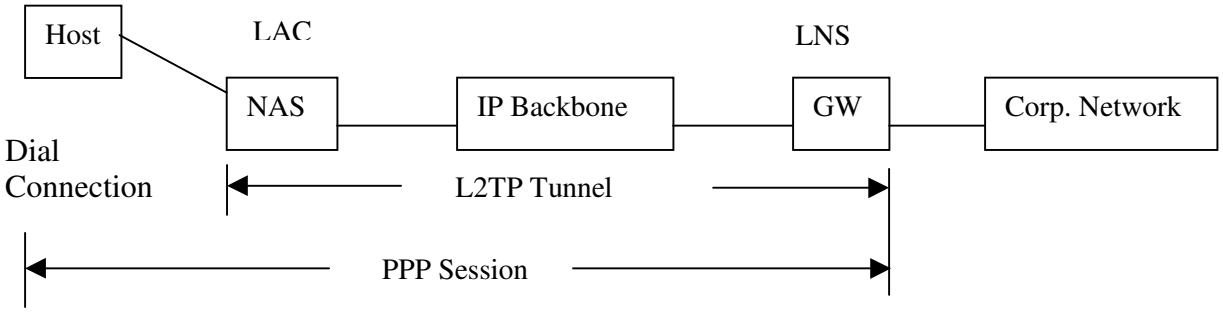

Figure 18 Compulsory Tunneling Example

5.1.2 Voluntary Tunneling

Voluntary tunneling refers to the case where an individual host/remote computer user connects to a remote site using a tunnel originating on the host, with no involvement from intermediate network nodes. There is more flexibility in how and to where they are created. This flexibility is particularly important for remote mobile users who may be dialing in from a new place each day. Because the ISP is not involved in the creation of the tunnel, the tunnel can span the networks of multiple ISPs without explicit configuration or agreement.

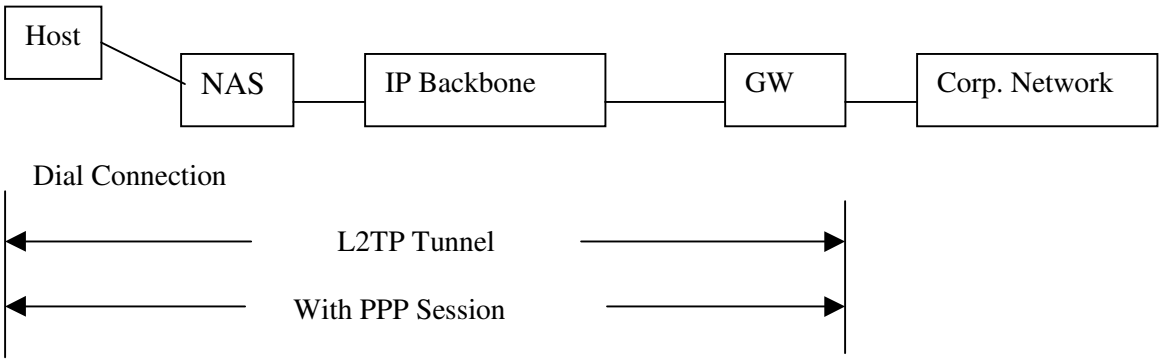

Figure 19 Voluntary Tunneling Example

The L2TP specification has support for voluntary tunneling, in so far as the LAC can be located on a host, not only on a network node. Note that such a host has two IP addresses – one for the LAC-LNS IP tunnel, and another allocated via PPP, for the network to which the host is connecting. The benefits of using L2TP for voluntary tunneling are that the existing authentication and address assignment mechanisms used by PPP can be reused without modification.

Because voluntary tunnels are initiated by the remote user and originate from the remote computer, there is more flexibility in how and to where they are created.

5.2 How does it work?

L2TP tunnels are initiated inside the service provider network & terminated on the customer premise. The central components of an L2TP networks are the LAC (Link Access Concentrator) & the LNS (L2TP Network Server). The LAC performs the following functions:

- Termination of modem & ISDN calls.
- First-level authentication and tunneling via RADIUS.
- Some level of initial PPP setup, much of which is overridden at a later stage by the LNS.
- Execution of the L2TP protocol in terms of command & control messages & encapsulation of the remote user's PPP traffic in L2TP packets.

The LNS can be thought of as that which resides "virtually" as a software function inside of another piece of networking equipment on the customer premises. It's responsible for executing key aspects of PPP, such as:

- Link Control Protocol (LCP)
- Network Control Protocol (NCP)

The LAC initiates a tunnel to the LNS using a series of L2TP command & control packets called "Attribute Value Pairs" or AVP's. AVP's are responsible for setting all tunnel parameters as well as the initialization, maintenance, and teardown of the tunnel.

As is the case with PPTP, L2TP features command and control packets as well as data packets. As it uses a UDP session, all packets are UDP encapsulated. Because of the connection-less aspect of L2TP, it utilizes more command & control messages than PPTP.

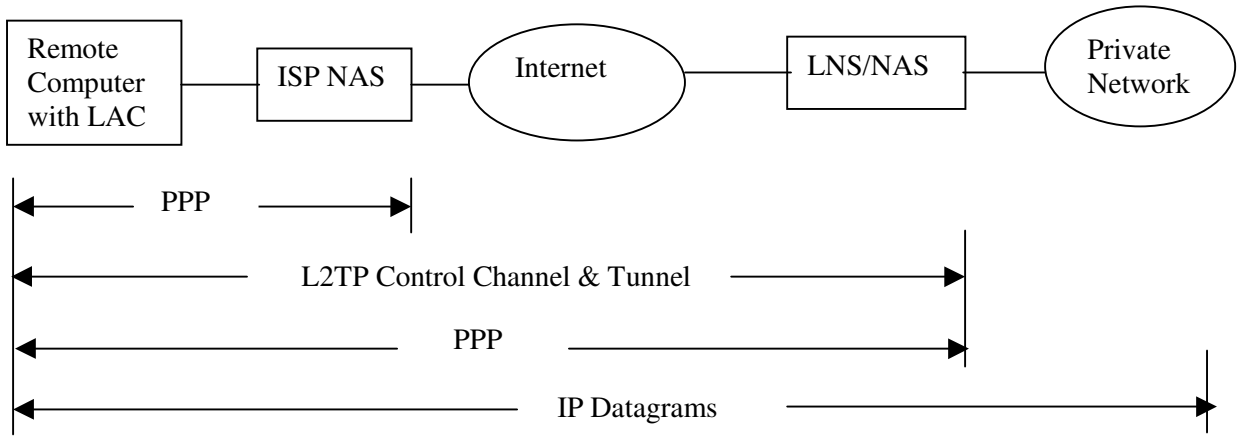

Figure 20 Remote user dial-in using L2TP

5.3 L2TP Protocol Characteristics

The characteristics of L2TP protocol are:

Multiplexing
L2TP has inherent support for the multiplexing of multiple calls from different users over a single link. Between the same two IP endpoints, there can be multiple L2TP tunnels, as identified by a tunnel-id, and multiple sessions within a tunnel, as identified by a session-id.

Signaling
This is supported via the in-built control connection protocol, allowing both tunnels and sessions to be established dynamically.

Data Security
By allowing for the transparent extension of PPP from the user, through the LAC to the LNS, L2TP allows for the use of whatever security mechanisms, with respect to both connection setup, and data transfer, may be used with normal PPP connections.

Multiprotocol Transport
L2TP transports PPP packets (and only PPP packets) and thus can be used to carry multiprotocol traffic since PPP itself is multiprotocol.

5.4 L2TP over Specific Media

L2TP is self-describing, operating at a level above the media over which it is carried. The following sections describe details needed to permit *interoperability* over specific media.

L2TP over UDP/IP

L2TP when using an UDP/IP media uses registered UDP port 1701. The entire L2TP packet, including payload and L2TP header, is sent within a UDP datagram. Once the source and destination addresses are established, they must remain until the tunnel is terminated. IP fragmentation may occur as the L2TP packet travels over the IP substrate. L2TP makes no special efforts to optimize this.
Also, it is proper to remember that the L2TP/UDP/IP transport is an unreliable transport.

IP

L2TP must offer the UDP encapsulation as described previously as its default configuration when operating in IP environs.

5.5 L2TP Security Considerations

L2TP suffers from the lack of solid tunnel protection mechanisms. Because L2TP encapsulates PPP, it inherits PPP's security mechanisms, including authentication and encryption services. PPP authenticates the client to the LNS but does not provide per-packet authentication. L2TP itself includes support for mutually authenticating the LAC and LNS tunnel endpoints at tunnel origination, but it too lacks stronger tunnel security mechanisms such as control and data packet protection. It doesn't provide key management facility, even though tunnel endpoint authentication relies on the distribution of tunnel passwords.

Tunnel Endpoint Security

The tunnel endpoints may optionally perform an authentication procedure of one another during tunnel establishment. This authentication has same security attributes as CHAP, and has reasonable protection against replay & snooping during the tunnel establishment process. This mechanism is not designed to provide authentication but just tunnel establishment only. So, it's quite easier for any intruder to snoop the tunnel stream to inject packets once an authenticated tunnel establishment has been completed successfully.
For authentication to occur, the LAC and LNS must share a single secret. Since a single secret is used and to guard against replay attacks, the tunnel authentication AVPs must include differentiating values in the CHAP ID fields for each message digest calculation.

Packet Level Security

The underlying transport makes available encryption, integrity and authentication services for all L2TP traffic for security. This secure transport operates on the entire L2TP packet and is functionally independent of PPP and the protocol being carried by PPP.

L2TP is concerned with confidentiality, authenticity, and integrity of the L2TP packets between its tunnel endpoints (the LAC and LNS).

End-to-End Security

Protecting the L2TP packet stream via a secure transport does, in turn, also protect the data within the tunneled PPP packets while transported from the LAC to the LNS.

So, summarizing the above said facts of L2TP as:

- Lacks tunnel protection mechanisms
- Lacks tunnel security mechanisms
- No key management facility but, authentication of tunnel endpoints relies on distribution of tunnel passwords

So, usually IPSec is used in conjunction with L2TP to protect L2TP traffic over IP and non-IP networks.

5.6 L2TP with IPSec

L2TP is a well-defined interoperable protocol that addresses the current shortcomings of IPSec-only client-to-gateway & gateway-to-gateway scenarios. By placing L2TP as payload within an IPSec packet, communications benefit from standards-based encryption & authenticity of IPSec.L2TP has broad vendor support particularly among the largest network access providers & has verified Interoperability. It helps accomplish user authentication, tunnel address assignment, multiprotocol support & multicast support using PPP. This combination is called L2TP/IPSec.

Due to incompatibilities between IKE protocol & Network Address translation, it is not possible to use L2TP/IPSec or IPSec tunnel mode through a NAT while taking advantage of automated key exchange.

The tunneling process is changed when using L2TP over IPSec. The L2TP data tunneling is performed through multiple levels of encapsulation. It's described as follows:

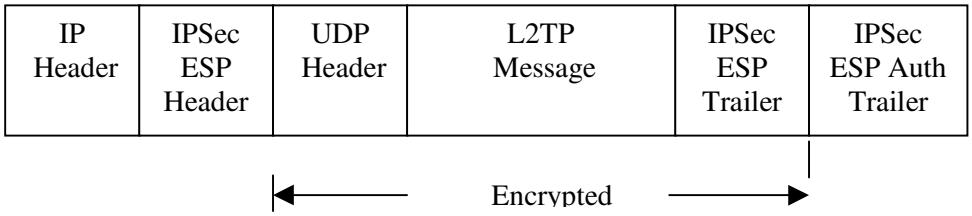

Figure 21 L2TP Encrypted control message

L2TP Encapsulation The initial PPP payload is encapsulated with a PPP header and an L2TP header.

UDP Encapsulation A UDP header is used to encapsulate the L2TP encapsulated packet with the source & destination ports set to 1701.

IPSec Encapsulation The encapsulated UDP header is then encrypted & encapsulated with an IPSec Encapsulating Security Payload (ESP) header & trailer & an IPSec Authentication (Auth) trailer.

IP Encapsulation The IPSec packet is encapsulated with a final IP header containing the source and destination IP addresses of the VPN client and VPN server.

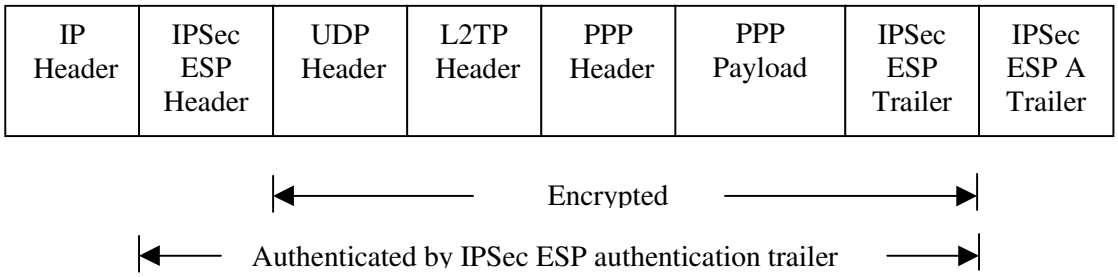

Figure 22 L2TP Tunneling Process

The L2TP working group recommends that L2TP command and control packets be secured by the use of IPSec Transport Mode. Figure 22 above shows an L2TP packet that has been secured using IPSec Transport Mode. The Unsecured area is “in the clear”, as it must be in order to traverse the IP backbone.

6 Comparison of protocols

	IPSec	PPTP	L2TP
OSI Layer	Layer 3	Layer 2	Layer 2

Table 2: Layers at which the three tunneling protocols are implemented.

Support for Communication Protocols:

	IPSec	PPTP	L2TP
Transported Protocols	IP	IP, IPX, NetBEUI	IP, IPX, NetBEUI
Required Underlying Protocol	IP	IP	IP, X.25, Frame Relay, ATM
Transport Layer Protocol	?	TCP for	
Number of tunnels supported	1	1	Several

Table 3: Support for Communication Protocols.

6.1 Security

Security is one of the basic requirements of VPN as the Internet almost does not provide any security guarantees, so the tunneling protocols should offer some security mechanisms. Of the above tunneling protocols, only IPSec provides the complete built-in security mechanisms. PPTP and L2TP do not provide data security functions but rely on PPP for their authentication and encryption services [9].

6.1.1 Authentication

Authentication is defined as the ability to verify parties on both ends of the link.
There are two levels of authentication- user authentication and packet authentication.
In user authentication, each party has to confirm the authenticity of the other party.
User authentication is not part of the IPSec protocol. PPTP and L2TP use the PPP authentication schemes (e.g., PAP, CHAP, and EAP) for providing user authentication.
In packet authentication, each packet carries some information that identifies the source of the packet to the receiver.
In IPSec, packet authentication is mainly provided by the AH header. Alternatively, ESP can also provide packet authentication. PPTP and L2TP do not have packet authentication scheme.

6.1.2 Integrity

The ability to verify that all data transmitted and received has not been tampered with or changed.

With IPSec integrity is managed by the AH and ESP headers. Both provide support for integrity, via the Integrity Check Value in the Authentication Data field of the header.
The PPTP packet assures integrity by tightly coordinating packet flow.
L2TP through the use of PPTP maintains integrity in the same manner.

6.1.3 Confidentiality

The ability to ensure that all transmitted data over the link is not read or intercepted by unauthorized users. The basis for maintaining confidentiality in all three protocols is through the use of encryption.
IPSec uses ESP to provide data confidentiality.
Encryption is provided by PPP based on a shared secret (the user's password).

6.1.4 Key Management

Key Management refers to managing the distribution of secret keys between the users of a network.
IPSec provides key management via Internet Key Exchange (IKE), whereas PPTP and L2TP do not provide any key management services.

Security Feature	IPSec	PPTP	L2TP
User Authentication/ Algorithm used	No	Yes/ PAP, CHAP, EAP, SPAP	Yes / PAP, CHAP, EAP, SPAP
Packet Authentication/ Algorithm used	Yes/HMAC SHA-1, HMAC MD5	No	No
Packet Encryption	Yes /DES, 3DES	No	No
Key Management	Yes /IKE	No	No

Table 4: Comparison of protocols based on security related features

6.1.5 Attacks on VPN

The possible attacks on a VPN are:

- Attacks against the protocol.
- Attack against the algorithms.
- Attacks against the implementation.
- Impersonation.
- Integrity.
- Disclosure
- Denial of service.

Attacks against the protocol

A cryptographic system is only as strong as the encryption algorithms and the hash functions it is based on. By breaking any of these, you break the whole system. The encryption algorithms and the key exchange protocols don't necessarily guarantee safety by themselves, and choosing an insecure seed for the key to be based on destroys the integrity of the protocol[10].

Attacks against the algorithm

The algorithm, or the mathematical operations that are performed on the data, could make the whole system weak. Proprietary encryption algorithms don't make a system weak. Many previously secret algorithms have been made public; reverse engineered, and, as a result, were shown to be ineffective. Using weak keys, using an insufficient amount of data size, and altering hash functions all contribute to the weakening of the system.

Attacks against the implementation

The vulnerability of key recovery is a major contributor to implementation attacks.
Some implementations leave temporary files, plaintext messages, and the data stored in buffers where they can easily be retrieved. Using a combination of keys is also a vulnerability (e.g.: Strong and weak key).

Impersonation - Impersonation attacks are those in which an attacker masquerades as another person. The strong authentication methods can reduce the effectiveness of impersonation attacks.

Integrity - Successful integrity attacks result in the undetected modification of the user data; for example, changing the contents of an electronic mail message in the transit. Integrity attacks are generally impossible to prevent. The best can be done is to detect the modification. Digital signatures of various types are useful defenses against integrity attacks.

Disclosure - Disclosure attacks result in the exposure of the data to an unintended person. The damage caused by disclosure attacks often depends upon the content of the data revealed: a meeting request may have little value to an opponent, whereas the disclosure of confidential sales projects could be ruinous. The typical defense against attacks is the use of strong encryption to hide network traffic.

Denial of service - Denial of service attacks are the hardest attacks to defend against, and the easiest to perpetrate. The purpose of these attacks, as the name suggests, is to deny service to valid users. The known denial of service attacks are teardrop, newtear, and syn flooding.

Common Attack on cryptographic Algorithms

Cipher text-only Attack - In a cipher text-only attack, the attacker knows nothing about the plaintext message, but given the cipher text, tries to make the guesses about the plaintext.

Known Plaintext Attack – In this attack, the attacker knows part of the plaintext document or can make an educated guess about it.

Chosen Plaintext Attack – In a chosen plaintext attack, the attacker takes some text and encrypts it with the unknown key.

Chosen Cipher text Attack – In a chosen cipher text attack, the attacker has the advantage of choosing an arbitrary selected cipher text and can find the corresponding plain text.

Man-in- the-Middle Attack – The "man-in-the-middle" hijacks the sender's and receiver's keys and substitutes his or her own, thereby giving the attacker the ability to intercept all future communication without either the sender or receiver knowing about it.

Timing Attack – This type of attack is relatively new and is based on measuring the execution times of a modular exponentiation operation that is used in cryptographic algorithms.

Brute-Force Attack – A brute-force attack is popular with attackers who have lot of computing power at their disposal. E.g. If f(x) = y, where y I is the cipher text, f(x) is the plaintext, and x is the key.

Differential Cryptanalysis – In differential cryptanalysis attacks, an attacker uses an iterative mapping process – that is, the mapping is based on a repeated function.

Internet Protocol Security (IPSec) Attacks

The Internet Security Protocol (IPSec) is not an encryption algorithm, and it is not an authentication algorithm. IPSec is a paradigm in which other algorithms protect data.

- Implementation Attacks.
- Key-Management Attacks.
- Key-Recovery/Export Law Attacks.
- Administrator and Wildcard Attacks.

Implementation Attacks - The IPSec standard only calls for one encryption algorithm (DES-CBC) and two authentication modes (HMAC-MD5 and HMAC-SHA-1); however it calls for the additional "NULL" algorithms, since AH and ESP may be optional. When a standard calls for an optional algorithm, it is trying to balance flexibility with security. The interpretation is that even if one end of the communication was to use DES-CBC, the other end should still be able to use the NULL, or no, algorithm and still communicate. The receiving end usually specifies the SA, but in order to be compatible with other systems, it must allow the NULL algorithm. Vendors could decide on how to implement this choice, thereby increasing the security exposure.

In the IPSec key-management protocol IKE component, both ends of the communications channel decide on how often the encryption keys should be changed. Given that many vendors support weaker, 40 –bit keys for backward compatibility, changing these keys now becomes critical, but still is a negotiated session. If it's is a weaker implementation, it's probably using longer time period, which in turn, gives an attacker more time to break the 40-bit key. Considering that 56-bit keys are now broken in three days, 40-bit keys shouldn't have even been part of the standard.

Key-Management Attacks - The protocol (IKE) specification specifies how these keys should be exchanged, but it usually refers to the start of the communication, not the end of it. There is a "time-out" mechanism in the public-key exchanges, and it was discovered that there isn't true *interoperability* between the vendors. In addition under the IKE specification, any side could terminate a session, but there is no way for the other end to know that the session has been terminated; the sending end would keep sending data. If

the station is sending data, what's to stop another station from receiving that data and, if the weak keys are used, spoofing the identity of the original host?

Key-Recovery/Export Law Attacks - There is really no such thing as a key-recovery/export law attack, but if an IPSec implementation is available in an international standard, it has one of two serious weaknesses. Either it will be IPSec using 40-bit keys (although, 56-bit IPSec was being released) or it will support key recovery.

Administrator and Wildcard Attacks - In IPSec there is a provision for an administrative and provision for wildcard matching. While there hasn't been a direct attack (at least none reported), some have argued that by even having an administrative interface to the SA (security association), you can potentially increase the chance that the interface can be attacked and the SA compromised. Since there is no provision for such as interface, it is left up to the vendor's implementation.

Point-to-point Tunneling Protocol (PPTP) Attacks

The PPTP protocol attack is an attack against the implementation**.**

- Attacking the GRE.
- Attacking the Passwords.

Attacking the GRE - PPP packets are encapsulated inside the GRE and tunneled via IP to their destination. GRE uses protocol number 47. GRE packets may carry a sequence number and an acknowledgement number and may use a sliding window to avoid congestion. This has some practical implications. It means that if we want to try and spoof the PPP packets encapsulated in GRE, we just need to desynchronize the GRE channel. This may be avoided by use of the sequence number; unfortunately, originally GRE didn't mandate the use of this sequence number, and it is therefore up to a vendor's particular implementation. The GRE didn't have a way for the end host to react to a bad or duplicate sequence number. It's possible that it can be just ignored, and then the PPP packets can be spoofed.

Attacking the Passwords - The PPTP authentication implementation supports three types of user authentication. The two that are concerned with security are the hashed method and the challenge response method. Hashed password authentication is based upon two one-way hashing functions. During the first hashing function, all passwords entered are converted to uppercase, which reduces the data space. Second, the hashing functions produce the same hash output, given the same password. Unfortunately there is no salt, so the hash outputs from the same input. Therefore in this authentication model, PPTP is open to dictionary attacks. In addition, both hash outputs are sent together in the communication string. An attacker can attack the first hash function to compromise the second hash function, thereby finding the password.

The second security authentication method uses the challenge Handshake Authentication Protocol (CHAP). CHAP works by the client contacting the server and the server sending back a challenge. The client then performs a hash function, adds some extra information, and sends this back to the server. The server looks in its own database and computes the hash with the challenge. If they are the same authentication succeeds. While this eliminates the dictionary attack, the hashing functions could still be attacked.

The PPTP framework calls for Microsoft's Point-to-Point Encryption (MPPE).The encryption is based on the user's password. After the initial communication is set up, only certain PPP packets

are encrypted. RFC-1700 lists those packets that are sent in the clear and those that are encrypted. MPPE then does not encrypt all the packets. This means you can attack the PPP protocol itself – for instance, spoofing the configuration packet containing certain DNS server information. MPPE uses RC4 cipher in either 40- or 128-bit key size. One of the main security problems lies in the fact that since there are no lowercase characters, a good selection of passwords from which to choose is eliminated. Therefore, claiming that PPTP is either 40-bit or 128-bit secure is incorrect. The session key is derived from the user's password. The password will have a much lower entropy. The only way one to reach true 40-bit or 128-bit entropy is by generating a random session key.

Moreover, since there is no salt, and since PPP uses common headers and trailers, it makes it a target for known-text types of attacks. In addition, since encryption is based on the user password, not a public key or shared secret key encryption algorithm authentication cannot be ensured.

Vulnerability with PPTP is that it relies on PPP. Prior to any communication, PPP sets up and initializes the communication parameters, and since PPP has no authentication against these packets, attacks like man-in-the-middle and spoofing may occur.

Attacks	IPSec	PPTP	L2TP
Denial-of-service attack	Resistant	Not Resistant	Not Resistant
Man-in-the-middle attack	Resistant	Not Resistant	Not Resistant
Dictionary attack	Resistant	Not Resistant	Not Resistant
Spoofing attack	Resistant	Not Resistant	Not Resistant

Table 5: Vulnerability of the protocols to security attacks

6.2 Performance

Performance is second to security when evaluating the tunneling protocols for VPN. Performance is measured in terms of throughput and latency.

Latency is essentially the communication delay, an expression of how much time it takes for a packet of data to get from one designated point to another. Encryption in a packet-based protocol further increases the latency, as the packets need to be reassembled in the correct order before decryption can occur.

Encapsulation requires adding information to each packet, which increases the packet size. This in turn increases the likelihood that internetwork routers will find the packets oversized and fragment them, further degrading performance. Packet fragmentation and data encryption can reduce dial-in system performance to unacceptable levels. Data compression can help solve this problem. However, the combination of compression and encapsulation requires additional computational power beyond that needed for security.

The cryptographic algorithm used also adds overhead. Cryptographic algorithm overhead is created by padding that must be added to packets for encryption and authentication algorithms before processing. The common encryption/decryption algorithms are block-based algorithms that operate on specific blocks of data. When data including minimum padding are not divisible by these block sizes, padding must be added to reach the desired block size prior to algorithmic processing.

PPTP uses UDP to carry the data packets and TCP to carry the command control packets. L2TP does not distinguish between packet types. All L2TP packets are UDP encapsulated.
As a connection-oriented protocol TCP requires an acknowledgment packet to come back for each chunk of data it sends out. A packet's TCP header is 12 bytes larger than a UDP Header. This results in degradation of throughput of PPTP as compared with L2TP that uses connectionless UDP.

PPTP may also have performance issues over high-latency networks. There are a couple of reasons for this. Here again, the first is the use of TCP for PPTP control packets. TCP is a session-oriented protocol, meaning a session exists between the PPTP client and the PPTP server during the lifecycle of the tunnel. TCP implements flow control based on configurable send and receive window sizes. The window size is the number of input or output buffers available for sending data. The size of the window is similar to the length of time a traffic light stays green before changing to yellow and red. When the road is wide open, the longer the light stays green the better (large window). Usually, larger window sizes lead to higher performance on faster networks. The problem is that performance over the Internet can fluctuate widely so it is difficult to predict an optimal window size.

L2TP utilizes more command and control messages than PPTP, because of the connection-less aspect of L2TP, but will also perform better over high latency networks.

6.3 Scalability

Scalability is measured in terms of the number of tunnels that a VPN gateway can sustain without degradation in performance. That is the number of simultaneous client-to-gateway as well as gateway-to-gateway connections depending on the topology.
IPSec and PPTP support just one tunnel between two users whereas L2TP supports multiple tunnels between users.

6.4 Flexibility

IPSec is quite flexible in terms of how encryption is implemented. It does not dictate an encryption algorithm- only the format of the encryption header itself.
As is the case with IPSec, there is more than one way to implement PPTP. Because it encapsulates the entire PPP frame, PPTP can support multiple network protocols (theoretically, any protocol that runs over PPP). We say theoretically because the PPTP server must have a forwarder for the appropriate protocol built into its system. PPTP is flexible in other ways as well. It can run over dial-up lines, local area networks (LANs), or wide area networks (WANs). It can run over the Internet or any other TCP/IP-based network, private or public.

6.5 Interoperability

Interoperability is the ability of a system or a product to work with other systems or products without special effort on the part of the customer.
PPTP is a vendor-specific, proprietary protocol, so interoperability is limited to products from supporting vendors. In contrast, L2TP is a multi-vendor effort, so interoperability is not as much of a problem.
Most currently available versions of IPSec are not interoperable between vendors. This reflects not only the status of the standards involved, but the options available for encryption, key management, and so on.

6.6 Multiprotocol support

IPSec only handles IP. It does not address network protocols other than IP. For non-IP traffic to be protected by IP, it first has to be tunneled inside of IP. This adds a layer of unnecessary overhead to data packets, so IPSec is not considered optimal when transporting protocols such as IPX/SPX, AppleTalk, or others.

6.7 Applications

Taking each protocol on its abilities and disregarding vendor-specific features, we found an appropriate use for each protocol.

Scenario	PPTP	L2TP	IPSec Transport Mode	Tunnel Mode
Remote Access VPN	√	√		
Branch Office VPN	√	√		√ (IP only)
Extranet using VPNs	√	√		√ (IP only)
Securing the Intranet			√ (Unicast only)	

Table 6: Comparison of protocols based on applications

7 Vendors

Vendor	Product	Type	Web Address	Technology
Checkpoint	Firewall –1	Firewall	http://www.checkpoint.com	IPSec
Trusted Information Systems	Gauntlet	Firewall	http://www.tis.com	IPSec
Raptor	Eagle Mobile NT 4.0	Firewall	http://www.raptor.com	IPSec
Cisco	Layer 2 Forwarding (L2F) Products	Hardware	http://www.cisco.com	L2F
Cisco	Cisco IOS	Software	http://www.cisco.com	IPSec, L2TP
3Com	3Access VPN products	Hardware	http://www.3com.com	IPSec
Compaq-Microcom	6000 Series	Hardware	http://www.microcom.com	PPTP
Extended Systems	Extended- Net VPN	Hardware	http://www.extendsys.com	PPTP with MPPE
RedCreek Communications	Ravlin	Hardware	http://www.redcreek.com	IPSec/ISAKMP/ Oakley
Timestep	Permit System	Hardware	http://www.timestep.com	IPSec
VPNet	VPN Remote client	Hardware	http://www.vpnet.com	IPSec
Information Resource Engineering	SafeNet/ Enterprise	Multiple	http://www.ire.com	IPSec
Novell	BorderManager	NOS	http://www.novell.com	PPTP
Microsoft	NT Server 4.0 Remote Access Server / PPTP	NOS	http://www.microsoft.com	PPTP and MPPE

8 Conclusion

Whenever there are competing protocols for a certain application, the most obvious question that arises is what protocol is best suited in a certain scenario? In the case of VPN tunneling protocols, there are three major protocols and a company or an individual trying to install a VPN solution would be faced with this question. Unfortunately, there is not one single "magic" answer to this question. The protocol selection depends on various factors. As a conclusion of our report, we would like to provide some pointers that can assist in this selection process.

The protocols running on the internal network of the company, looking for a VPN solution, will play a major role in this decision. This is because IPsec and other VPN protocols are not compatible with every other protocol. The bottom line here is that if your network is running only on TCP/IP, you have the common denominator of all VPN technologies. Every hardware and software VPN solution, no matter the protocol, is designed to package and tunnel TCP/IP packets. NetBEUI or IPX/SPX-based networks will have a limited number of options if your VPN solution requires those packets to be tunneled to another site, in this case, your best bet would be to go with either PPTP or L2TP.

When setting up a VPN network that must work through a NAT box and security is not an issue, PPTP would be the most suitable tunneling protocol. PPTP will suit most client-to-gateway or gateway-to-gateway scenarios and offer simple User Authentication via CHAP or PAP. PPTP or L2TP offer multiprotocol support, but remember that L2TP on it's own lacks any kind of data confidentiality (encryption). Both L2TP and PPTP also lack Machine Authentication and must rely on User Authentication.

Of all the VPN protocols, IPSec offers the strongest security (notably encryption, authentication, and key management); however, IPSec doesn't feature User Authentication or any multiprotocol support. Many vendors implemented solutions to resolve problems like User Authentication by adding RADIUS Authentication to their product offerings. IPSec is also extremely flexible to an extent where configuration is an issue, but it offers the user the best security mechanisms available and is suitable for almost any gateway-to-gateway scenario.

A single VPN protocol cannot fulfil every customer requirement, as there will always be trade-offs between functionality, interoperability and security. The most comprehensive VPN solution to date is a combination of two protocols - *L2TP with IPSec*. This hybrid protocol will offer a robust, very secure mechanism for creating both Client-to-LAN and LAN-to-LAN VPNs that are easy to set up and use as PPTP. This combination protocol also offers maximum interoperability.

Consider PPTP if you have low security requirements, need a simple VPN solution and multi-protocol support is a must.

Consider L2TP if you need a faster and leaner solution than offered by PPTP.

Consider IPSec if the main selection criterion is security and you need ease of use and configuration.

Consider L2TP/IPSec if complete interoperability and strong security are most important to you.

References

[1] Kent, S. and Atkinson, R., Security Architecture for the Internet Protocol", RFC 2401, November 1998.

[2] Townsley, W. et al, "Layer Two Tunneling Protocol - L2TP", RFC 2661, August 1999.

[3] Gleeson, B. et al, "A Framework for IP Based Virtual Private Networks", RFC 2764, February 2000.

[4] Hamzeh, K. et al, "Point-to-Point Tunneling Protocol (PPTP)", RFC 2637, July 1999.

[5] Oppliger, R., "Security Technologies for the World Wide Web", Artech House Computer Library, 2000.

[6] Venkateswaran, R., "Virtual Private Networks", IEEE Potentials Magazine, February/March 2001.

[7] Yuan, R. and Strayer, T., "Virtual Private Networks, Addison-Wesley, 2001.

[8] Perlmutter, B. and Zarkover, J., "Virtual Private Networking: A View from the Trenches", Prentice Hall PTR, 2000.

[9] 3Com White Paper, "Virtual Private Networks: Internet-based VPNs", February 2001, http://www.3com.com

[10] Brown, S., "Implementing Virtual Private Networks", McGraw Hill, 1999.

[11] Microsoft White Paper, "Microsoft Privacy Protected Network Access: Virtual Private Networking and Intranet Security", May 1999, http://www.microsoft.com