

Secure Remote Access with L2TP

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2000). All Rights Reserved.

Abstract

L2TP protocol is a virtual extension of PPP across IP network infrastructure. L2TP makes possible for an access concentrator (LAC) to be near remote clients, while allowing PPP termination server (LNS) to be located in enterprise premises. L2TP allows an enterprise to retain control of RADIUS data base, which is used to control Authentication, Authorization and Accountability (AAA) of dial-in users. The objective of this document is to extend security characteristics of IPsec to remote access users, as they dial-in through the Internet. This is accomplished without creating new protocols and using the existing practices of Remote Access and IPsec. Specifically, the document proposes three new RADIUS parameters for use by the LNS node, acting as Secure Remote Access Server (SRAS) to mandate network level security between remote clients and the enterprise. The document also discusses limitations of the approach.

1. Introduction and Overview

Now-a-days, it is common practice for employees to dial-in to their enterprise over the PSTN (Public Switched Telephone Network) and perform day-to-day operations just as they would if they were in corporate premises. This includes people who dial-in from their home and road warriors, who cannot be at the corporate premises. As the Internet has become ubiquitous, it is appealing to dial-in through the Internet to save on phone charges and save the dedicated voice lines from being clogged with data traffic.

The document suggests an approach by which remote access over the Internet could become a reality. The approach is founded on the well-known techniques and protocols already in place. Remote Access extensions based on L2TP, when combined with the security offered by IPsec can make remote access over the Internet a reality. The approach does not require inventing new protocol(s).

The trust model of remote access discussed in this document is viewed principally from the perspective of an enterprise into which remote access clients dial-in. A remote access client may or may not want to enforce end-to-end IPsec from his/her end to the enterprise. However, it is in the interest of the enterprise to mandate security of every packet that it accepts from the Internet into the enterprise. Independently, remote users may also pursue end-to-end IPsec, if they choose to do so. That would be in addition to the security requirement imposed by the enterprise edge device.

Section 2 has reference to the terminology used throughout the document. Also mentioned are the limited scope in which some of these terms may be used in this document. Section 3 has a brief description of what constitutes remote access. Section 4 describes what constitutes network security from an enterprise perspective. Section 5 describes the model of secure remote access as a viable solution to enterprises. The solution presented in section 5 has some limitations. These limitations are listed in section 6. Section 7 is devoted to describing new RADIUS attributes that may be configured to turn a NAS device into Secure Remote Access Server.

2. Terminology and scope

Definition of terms used in this document may be found in one of (a) L2TP Protocol document [Ref 1], (b) IP security Architecture document [Ref 5], or (c) Internet Key Exchange (IKE) document [Ref 8].

Note, the terms Network Access Server (NAS) and Remote Access Server(RAS) are used interchangeably throughout the document. While PPP may be used to carry a variety of network layer packets, the focus of this document is limited to carrying IP datagrams only.

"Secure Remote Access Server" (SRAS) defined in this document refers to a NAS that supports tunnel-mode IPsec with its remote clients. Specifically, LNS is the NAS that is referred. Further, involuntary tunneling is assumed for L2TP tunnel setup, in that remote clients initiating PPP session and the LAC that tunnels the PPP sessions are presumed to be distinct physical entities.

Lastly, there are a variety of transport mediums by which to tunnel PPP packets between a LAC and LNS. Examples include Frame Relay or ATM cloud and IP network infrastructure. For simplicity, the document assumes a public IP infrastructure as the medium to transport PPP packets between LAC and LNS. Security of IP packets (embedded within PPP) in a trusted private transport medium is less of a concern for the purposes of this document.

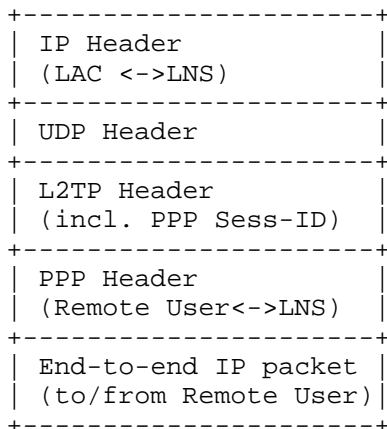
3. Remote Access operation

Remote access is more than mere authentication of remote clients by a Network Access Server(NAS). Authentication, Authorization, Accounting and routing are integral to remote access. A client must first pass the authentication test before being granted link access to the network. Network level services (such as IP) are granted based on the authorization characteristics specified for the user in RADIUS. Network Access Servers use RADIUS to scale for large numbers of users supported. NAS also monitors the link status of the remote access clients.

There are a variety of techniques by which remote access users are connected to their enterprise and the Internet. At a link level, the access techniques include ISDN digital lines, analog plain-old-telephone-service lines, xDSL lines, cable and wireless to name a few. PPP is the most common Layer-2 (L2)protocol used for carrying network layer packets over these remote access links. PPP may be used to carry a variety of network layer datagrams including IP, IPX and AppleTalk. The focus of this document is however limited to IP datagrams only.

L2TP is a logical extension of PPP over an IP infrastructure. While a LAC provides termination of Layer 2 links, LNS provides the logical termination of PPP. As a result, LNS becomes the focal point for (a) performing the AAA operations for the remote users, (b) assigning IP address and monitoring the logical link status (i.e., the status of LAC-to-LNS tunnel and the link between remote user and LAC), and (c) maintaining host-route to remote user network and providing routing infrastructure into the enterprise.

L2TP uses control messages to establish, terminate and monitor the status of the logical PPP sessions (from remote user to LNS). These are independent of the data messages. L2TP data messages contain an L2TP header, followed by PPP packets. The L2TP header identifies the PPP session (amongst other things) to which the PPP packet belongs. The IP packets exchanged from/to the remote user are carried within the PPP packets. The L2TP data messages, carrying end-to-end IP packets in an IP transport medium may be described as follows. The exact details of L2TP protocol may be found in [Ref 1].



4. Requirements of an enterprise Security Gateway

Today's enterprises are aware of the various benefits of connecting to the Internet. Internet is a vast source of Information and a means to disseminate information and make available certain resources to the external world. However, enterprises are also aware that security breaches (by being connected to the Internet) can severely jeopardize internal network.

As a result, most enterprises restrict access to a pre-defined set of resources for external users. Typically, enterprises employ a firewall to restrict access to internal resources and place externally accessible servers in the DeMilitarized Zone (DMZ), in front of the firewall, as described below in Figure 1.

becomes very difficult to allow employees to access corporate resources, via the Internet, even if you are willing to forego security over the Internet.

With the advent of IPsec, it is possible to secure corporate data across the Internet by employing a Security Gateway within the enterprise. Firewall may be configured to allow IKE and IPsec packets directed to a specific Security Gateway behind the firewall. It then becomes the responsibility of the Security Gateway to employ the right access list for external connections seeking entry into the enterprise. Essentially, the access control functionality for IPsec secure packets would be shifted to the Security Gateway (while the access control for clear packets is retained with the firewall). The following figure illustrates the model where a combination of Firewall and Security Gateway control access to internal resources.

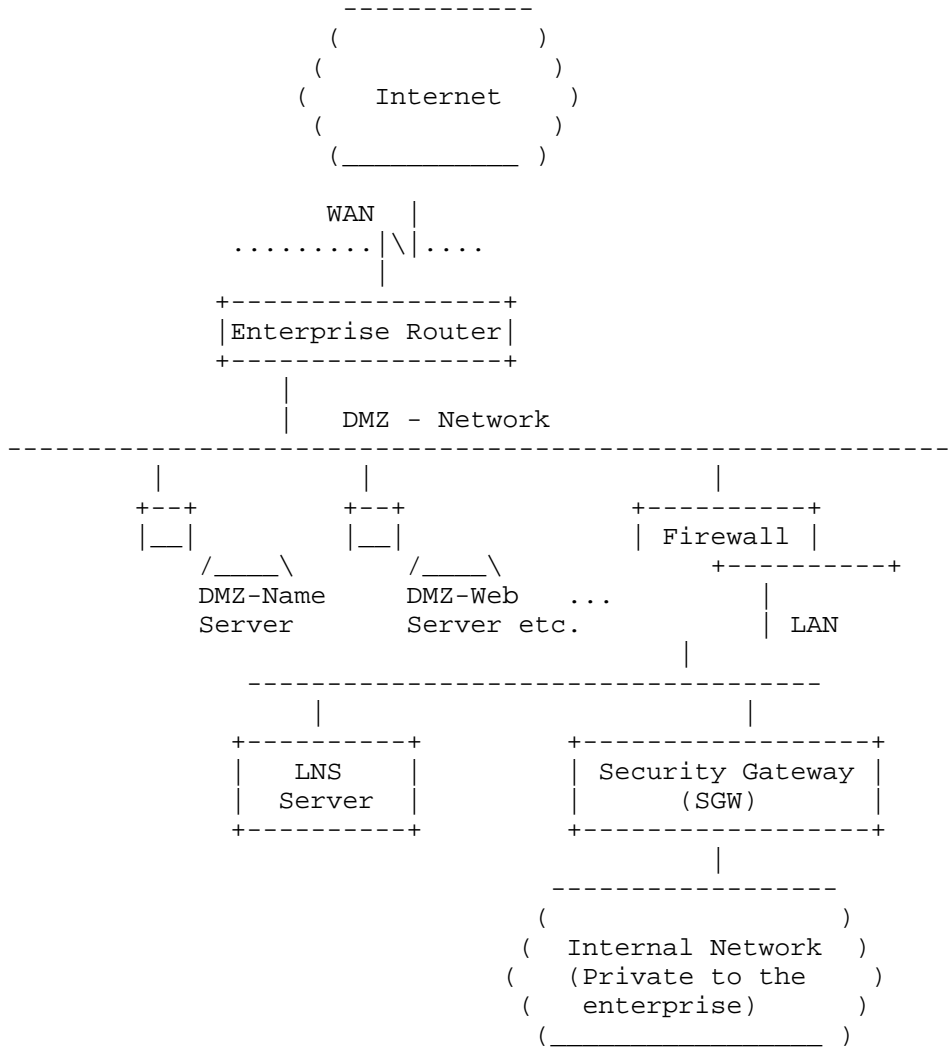


Figure 2: Security Model based on Firewall and Security Gateway

In order to allow employee dial-in over the Internet, an LNS may be placed behind a firewall, and the firewall may be configured to allow UDP access to the LNS from the Internet. Note, it may not be possible to know all the IP addresses of the LACs located on the Internet at configuration time. Hence, the need to allow UDP access from any node on the Internet. The LNS may be configured to process only the L2TP packets and drop any UDP packets that are not L2TP.

Such a configuration allows remote access over the Internet. However, the above setup is prone to a variety of security attacks over the Internet. It is easy for someone on the Internet to steal a remote access session and gain access to precious resources of the enterprise. Hence it is important that all packets are preserved with IPsec to a security Gateway (SGW) behind the LNS, so the Security Gateway will not allow IP packets into corporate network unless it can authenticate the same.

The trust model of secure remote access assumes that the enterprise and the end user are trusted domains. Everything in between is not trusted. Any examination of the end-to-end packets by the nodes enroute would violate this trust model. From this perspective, even the LAC node enroute must not be trusted with the end-to-end IP packets. Hence, location and operation of LAC is not relevant for the discussion on security. On the other hand, location and operation of LNS and the Security Gateway (SGW) are precisely the basis for discussion.

Having security processing done on an independent Security gateway has the following shortcomings.

1. Given the trust model for remote access, the SGW must be configured with a set of security profiles, access control lists and IKE authentication parameters for each user. This mandates an independent provisioning of security parameters on a per-user basis. This may not be able to take advantage of the user-centric provisioning on RADIUS, used by the LNS node.
2. Unlike the LNS, SGW may not be in the routing path of remote access packets. I.e., there is no guarantee that the egress IP packets will go through the chain of SGW and LNS before they are delivered to remote user. As a result, packets may be subject to IPSec in one direction, but not in the other. This can be a significant threat to the remote access trust model.
3. Lastly, the SGW node does not have a way to know when a remote user node(s) simply died or the LAC-LNS tunnel failed. Being unable to delete the SAs for users that no longer exist could drain the resources of the SGW. Further, the LNS cannot even communicate the user going away to the SGW because, the SGW maintains its peer nodes based on IKE user ID, which could be different the user IDs employed by the LNS node.

5. Secure Remote Access

Combining the functions of IPsec Security Gateway and LNS into a single system promises to offer a viable solution for secure remote access. By doing this, remote access clients will use a single node as both (a) PPP termination point providing NAS service, and (b) the Security gateway node into the enterprise. We will refer this node as "Secure Remote Access Server" (SRAS).

The SRAS can benefit greatly from the confluence of PPP session and IPsec tunnel end points. PPP session monitoring capability of L2TP directly translates to being able to monitor IPsec tunnels. Radius based user authorization ability could be used to configure the security characteristics for IPsec tunnel. This includes setting access control filters and security preferences specific to each user. This may also be extended to configuring IKE authentication and other negotiation parameters, when automated key exchange is solicited. Security attributes that may be defined in Radius are discussed in detail in section 7. Needless to say, the centralized provisioning capability and scalability of Radius helps in the configuration of IPsec.

As for remote access, the benefit is one of IPsec security as befitting the trust model solicited by enterprises for the end-to-end IP packets traversing the Internet. You may use simply AH where there is no fear of external eaves-dropping, but you simply need to authenticate packet data, including the source of packet. You may use ESP (including ESP-authentication), where there is no trust of the network and you do not want to permit eaves-dropping on corporate activities.

Operation of SRAS requires that the firewall be configured to permit UDP traffic into the SRAS node. The SRAS node in turn will process just the L2TP packets and drop the rest. Further, the SRAS will require all IP packets embedded within PPP to be one of AH and ESP packets, directed to itself. In addition, the SRAS will also permit IKE UDP packets (with source and destination ports sets to 500) directed to itself in order to perform IKE negotiation and generate IPsec keys dynamically. All other IP packets embedded within PPP will be dropped. This enforces the security policy for the enterprise by permitting only the secure remote access packets into the enterprise. When a PPP session is dropped, the IPsec and ISAKMP SAs associated with the remote access user are dropped from the SRAS. All the shortcomings listed in the previous section with LNS and SGW on two systems disappear with the Secure Remote Access Server. Figure 3 below is a typical description of an enterprise supporting remote access users using SRAS system.

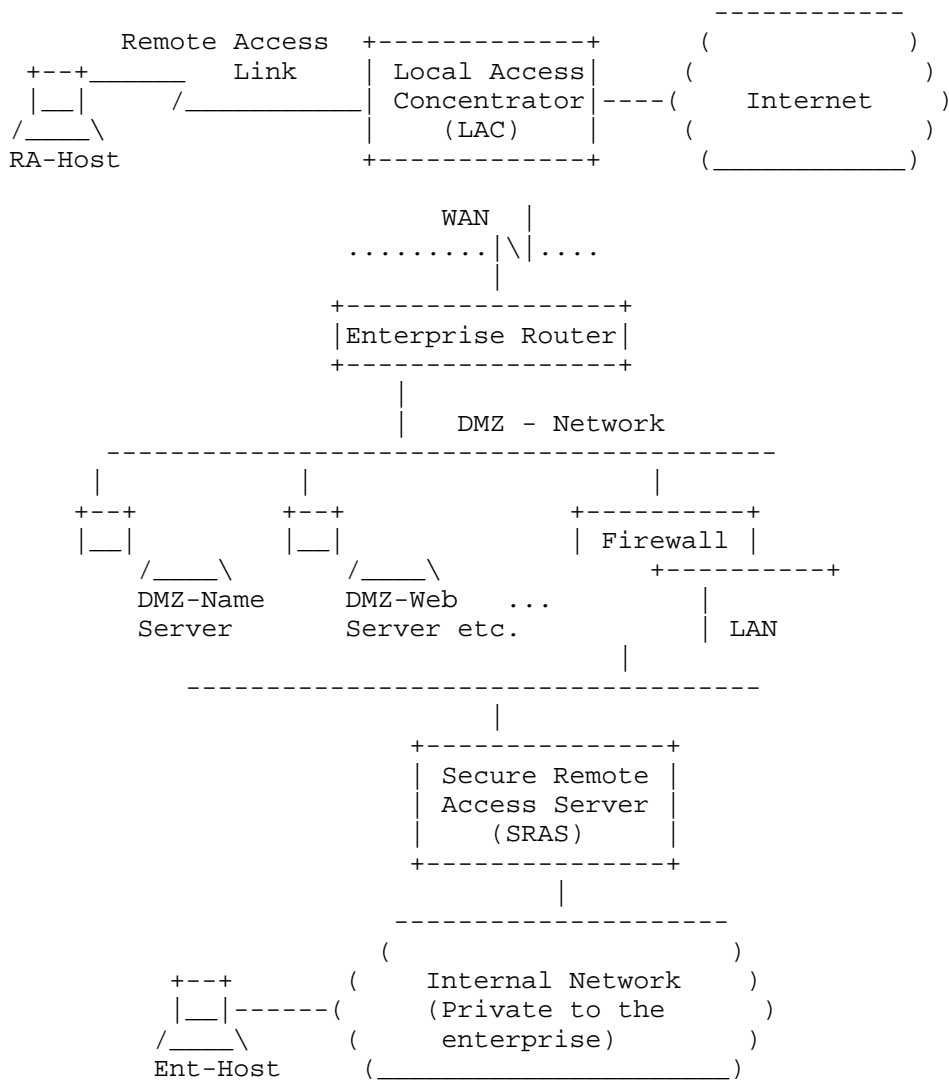
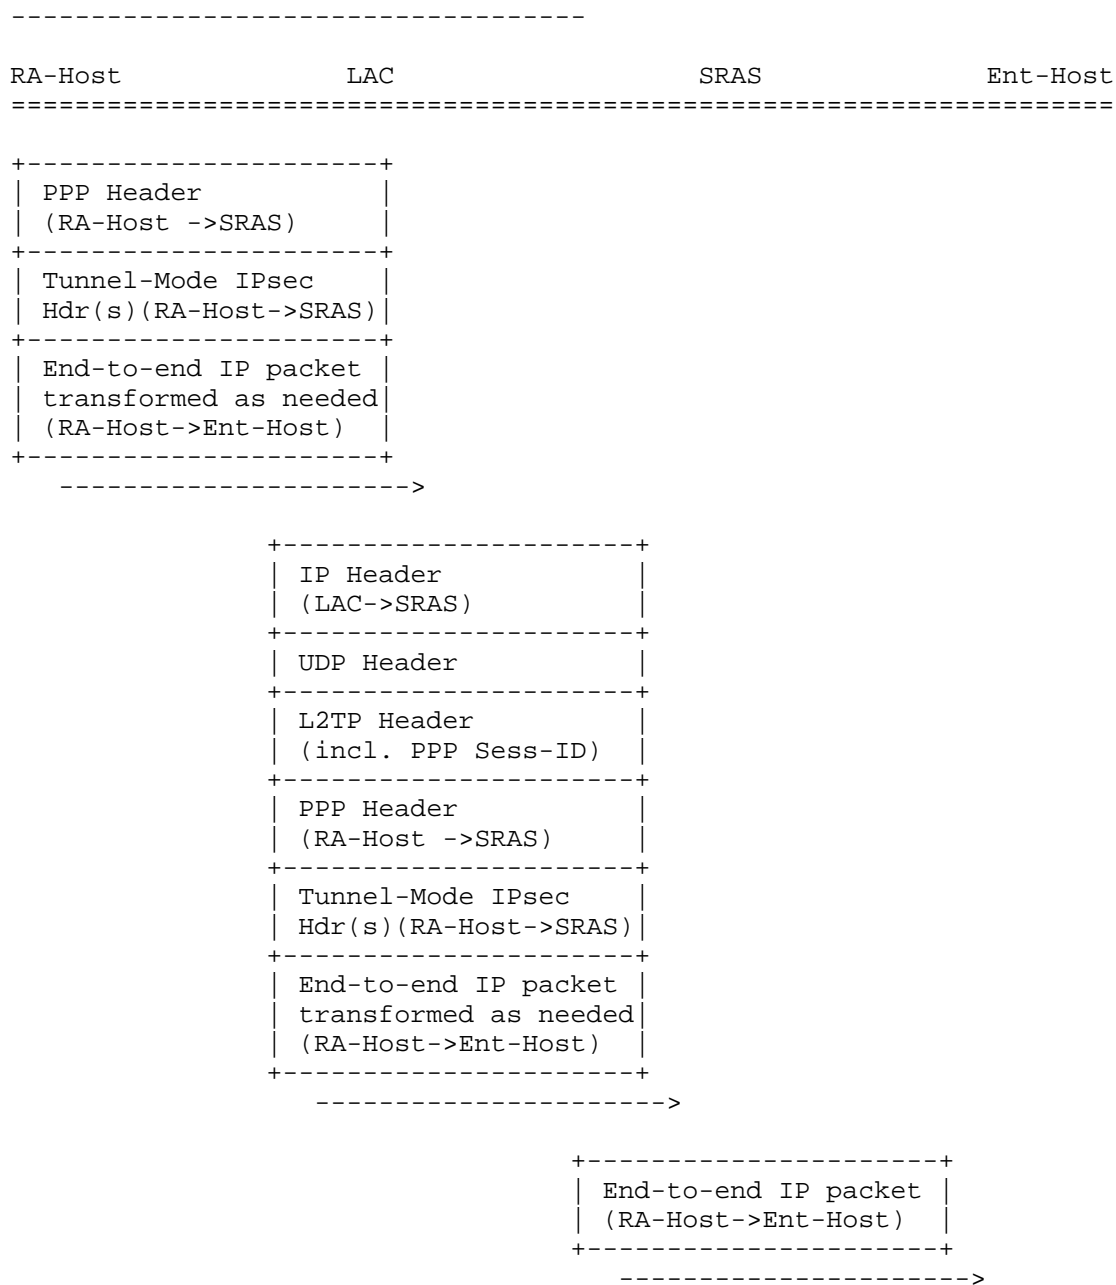


Figure 3: Secure Remote Access Server operation in an Enterprise

The following is an illustration of secure remote access data flow as end-to-end IP packets traverse the Internet and the SRAS. The example shows IP packet tunneling and IPsec transformation as packets are exchanged between a remote Access host (RA-Host) and a host within the enterprise (say, Ent-Host).

Note, the IP packets originating from or directed to RA-Host are shown within PPP encapsulation, whereas, all other packets are shown simply as IP packets. It is done this way to highlight the PPP packets encapsulated within L2TP tunnel. The PPP headers below are identified by their logical source and destination in parenthesis. Note, however, the source and recipient information of the PPP data is not a part of PPP header. This is described thus, just for clarity. In the case of an L2TP tunnel, the L2TP header carries the PPP session ID, which indirectly identifies the PPP end points to the LAC and the LNS. Lastly, the IPsec Headers section below include the tunneling overhead and the AH/ESP headers that are attached to the tunnel.

RA-Host to Ent-Host Packet traversal:



Ent-Host to RA-Host Packet traversal:

```

Ent-Host          SRAS          LAC          RA-Host
=====

```

```

+-----+
| End-to-end IP packet |
| (Ent-Host->Ra-Host) |
+-----+
----->

```

```

+-----+
| IP Header             |
| (SRAS->LAC)           |
+-----+
| UDP Header            |
+-----+
| L2TP Header           |
| (incl. PPP Sess-ID)  |
+-----+
| PPP Header            |
| (SRAS->RA-Host)      |
+-----+
| Tunnel-Mode IPsec    |
| Hdr(s) (SRAS->RA-Host)|
+-----+
| End-to-end IP packet |
| transformed as needed|
| (Ent-Host->RA-Host)  |
+-----+
----->

```

```

+-----+
| PPP Header            |
| (SRAS->RA-Host)      |
+-----+
| Tunnel-Mode IPsec    |
| Hdr(s) (SRAS->RA-Host)|
+-----+
| End-to-end IP packet |
| transformed as needed|
| (Ent-Host->RA-Host)  |
+-----+
----->

```

6. Limitations to Secure Remote Access using L2TP

The SRAS model described is not without its limitations. Below is a list of the limitations.

1. Tunneling overhead: There is considerable tunneling overhead on the end-to-end IP packet. Arguably, there is overlap of information between tunneling headers. This overhead will undercut packet throughput.

The overhead is particularly apparent at the LAC and SRAS nodes. Specifically, the SRAS has the additional computational overhead of IPsec processing on all IP packets exchanged with remote users. This can be a significant bottleneck in the ability of SRAS to scale for large numbers of remote users.

2. Fragmentation and reassembly: Large IP packets may be required to undergo fragmentation and reassembly at the LAC or the LNS as a result of multiple tunnel overhead tagged to the packet. Fragmentation and reassembly can have an impact on packet throughput and latency. However, it is possible to avoid the overhead by reducing the MTU permitted within PPP frames.
3. Multiple identity and authentication requirement: Remote Access users are required to authenticate themselves to the SRAS in order to be obtain access to the link. Further, when they require the use of IKE to automate IPsec key exchange, they will need to authenticate once again with the same or different ID and a distinct authentication approach. The authentication requirements of IKE phase 1 [Ref 8] and LCP [Ref 3] are different.

However, it is possible to have a single authentication approach (i.e., a single ID and authentication mechanism) that can be shared between LCP and IKE phase 1. The Extended Authentication Protocol(EAP) [Ref 4] may be used as the base to transport IKE authentication mechanism into PPP. Note, the configuration overhead is not a drag on the functionality per se.

4. Weak security of Link level authentication: As LCP packets traverse the Internet, the Identity of the remote user and the password (if a password is used) is sent in the clear. This makes it a target for someone on the net to steal the information and masquerade as remote user. Note, however, this type of password stealing will not jeopardize the security of the enterprise per se, but could result in denial of service to remote users. An intruder can collect the password data and simply steal the link, but will not be able to run any IP applications subsequently, as the SRAS will fail non-IPsec packet data.

A better approach would be to employ Extended Authentication Protocol (EAP) [Ref 4] and select an authentication technique that is not prone to stealing over the Internet. Alternately, the LAC and the SRAS may be independently configured to use IPsec to secure all LCP traffic exchanged between themselves.

7. Configuring RADIUS to support Secure Remote Access.

A centralized RADIUS database is used by enterprises to maintain the authentication and authorization requirements of the dial-in Users. It is also believed that direct dial-in access (e.g., through the PSTN network is) safe and trusted and does not need any scrutiny outside of the link level authentication enforced in LCP. This belief is certainly not shared with the dial-in access through the Internet.

So, while the same RADIUS database may be used for a user directly dialing-in or dialing in through the Internet, the security requirements may vary. The following RADIUS attributes may be used to mandate IPsec for the users dialing-in through the Internet. The exact values for the attributes and its values may be obtained from IANA (refer Section 10).

7.1. Security mandate based on access method

A new RADIUS attribute IPSEC_MANDATE (91) may be defined for each user. This attribute may be given one of the following values.

NONE	(=0)	No IPsec mandated on the IP packets embedded within PPP.
LNS_AS_SRAS	(=1)	Mandates Tunnel mode IPsec on the IP packets embedded within PPP, only so long as the PPP session terminates at an LNS. LNS would be the tunnel mode IPsec end point.
SRAS	(=2)	Mandates Tunnel mode IPsec on the IP packets embedded within PPP, irrespective of the NAS type the PPP terminates in. I.e., the IPsec mandate is not specific to LNS alone, and is applicable to any NAS, terminating PPP. NAS would be the tunnel mode IPsec end point.

When IPSEC_MANDATE attribute is set to one of LNS_AS_SRAS or SRAS, that would direct the NAS to drop any IP packets in PPP that are not associated with an AH or ESP protocol. As an exception, the NAS will continue to process IKE packets (UDP packets, with source and destination port set to 500) directed from remote users. Further, the security profile parameter, defined in the following section may add additional criteria for which security is not mandatory.

7.2. Security profile for the user

A new SECURITY_PROFILE (92) parameter may be defined in RADIUS to describe security access requirements for the users. The profile could contain information such as the access control security filters, security preferences and the nature of Keys (manual or automatic generated via the IKE protocol) used for security purposes.

The SECURITY-PROFILE attribute can be assigned a filename, as a string of characters. The contents of the file could be vendor specific. But, the contents should include (a) a prioritized list access control security policies, (b) Security Association security preferences associated with each security policy.

7.3. IKE negotiation profile for the user

If the security profile of a user requires dynamic generation of security keys, the parameters necessary for IKE negotiation may be configured separately using a new IKE_NEGOTIATION_PROFILE (93) parameter in RADIUS. IKE-NEGOTIATION_PROFILE attribute may be assigned a filename, as a string of characters. The contents of the file could however be vendor specific. The contents would typically include (a) the IKE ID of the user and SRAS, (b) preferred authentication approach and the associated parameters, such as a pre-shared-key or a pointer to X.509 digital Certificate, and, (c) ISAKMP security negotiation preferences for phase I.

8. Acknowledgements

The author would like to express sincere thanks to Steve Willens for initially suggesting this idea. The author is also thankful to Steve for the many informal conversations which were instrumental in the author being able to appreciate the diverse needs of the Remote Access area.

9. Security Considerations

This document is about providing secure remote access to enterprises via the Internet. However, the document does not address security issues for network layers other than IP. While the document focus is on security over the Internet, the security model provided is not limited to the Internet or the IP infrastructure alone. It may also be applied over other transport media such as Frame Relay and ATM clouds. If the transport media is a trusted private network infrastructure, the security measures described may not be as much of an issue. The solution suggested in the document is keeping in view the trust model between a remote user and enterprise.

10. IANA Considerations

This document proposes a total of three new RADIUS attributes to be maintained by the IANA. These attributes IPSEC_MANDATE, SECURITY_PROFILE and IKE_NEGOTIATION_PROFILE may be assigned the values 91, 92 and 93 respectively so as not to conflict with the definitions for recognized radius types, as defined in <http://www.isi.edu/in-notes/iana/assignments/radius-types>.

The following sub-section explains the criteria to be used by the IANA to assign additional numbers as values to the IPSEC-MANDATE attribute described in section 7.1.

10.1. IPSEC-MANDATE attribute Value

Values 0-2 of the IPSEC-MANDATE-Type Attribute are defined in Section 7.1; the remaining values [3-255] are available for assignment by the IANA with IETF Consensus [Ref 11].

REFERENCES

- [1] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G. and B. Palter, "Layer Two Tunneling Protocol L2TP", RFC 2661, August 1999.
- [2] Rigney, C., Rubens, A., Simpson, W. and S. Willens, "Remote Authentication Dial In User Service (RADIUS)", RFC 2138, April 1997.
- [3] Simpson, W., "The Point-to-Point Protocol (PPP)", STD 51, RFC 1661, July 1994.
- [4] Blunk, L. and Vollbrecht, J. "PPP Extensible Authentication Protocol (EAP)", RFC 2284, March 1998.

- [5] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.
- [6] Kent, S. and R. Atkinson, "IP Encapsulating Security Payload (ESP)", RFC 2406, November 1998.
- [7] Kent, S. and R. Atkinson, "IP Authentication Header", RFC 2402, November 1998.
- [8] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, November 1998.
- [9] Piper, D., "The Internet IP Security Domain of Interpretation for ISAKMP", RFC 2407, November 1998.
- [10] Reynolds, J. and J. Postel, "Assigned Numbers", STD 2, RFC 1700, October 1994.
See also <http://www.iana.org/numbers.html>
- [11] Narten, T. and H. Alvestrand, "Guidelines for writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.
- [12] Meyer, G., "The PPP Encryption Control Protocol (ECP)", RFC 1968, June 1996.
- [13] Sklower, K. and G. Meyer, "The PPP DES Encryption Protocol, Version 2 (DESE-bis)", RFC 2419, September 1998.

Author's Address

Pyda Srisuresh
Campio Communications
630 Alder Drive
Milpitas, CA 95035
U.S.A.

Phone: +1 (408) 519-3849
EMail: srisuresh@yahoo.com

Full Copyright Statement

Copyright (C) The Internet Society (2000). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.