INTERNET-DRAFT

George Tsirtsis, BTLABS
Alan O'Neill, BTLABS
January 1998

NAT Bypass for End 2 End 'sensitive' applications

<draft-tsirtsis-nat-bypass-00.txt>

Status of this Memo

Abstract

This document attempts to generate discussion on methods to run end 2 end 'sensitive' protocols and capabilities, like IPSEC, on networks that use Network Address Translators (NAT). The proposal does so by outlining one method to bypass NAT, when the required capabilities cannot be supported by NAT. The method uses a tunnel between a local host and the NAT box in order to dynamically allocate addresses to those hosts that need to communicate with external networks. With an allocated external address, the local hosts are able to communicate with external hosts without breaking the end 2 end principle. This proposal does not introduce any new protocols, it simply reuses existing protocols to provide an example solution.

Terminology

Application

In this document the term application refers to anything that uses the IP network protocol (IPSEC, FTP etc).

L2TP

In this document only the end 2 end flavour of L2TP is considered

(otherwise known as PPTP) were the Host and the L2TP Access Concentrator (LAC) are both in the Host and only the connection between LAC and L2TP Network Server (LNS) goes across the network.

Bypass the NAT

This means that the Address Translation function is bypassed, NOT the NAT box, since the tunnel that bypasses the translation function, MAY terminate at a Router combined with a NAT box.

1. Introduction and Motivation

Network Address Translation (NAT) is used today as an interim solution to the problem of limited address space in IPv4. One can design a network using private address space (not globally unique)and use NAT in order to allow communication with external networks. The NAT typically has only a small number of external addresses available, resulting in savings in the IPv4 address space.

Unfortunately, address translation breaks one of the fundamental principles of Internet; the End 2 End Principle [ROUT]. This recommends that packets flow end to end, between hosts, without anyone changing its contents along the path. A number of applications have been designed with that principle in mind and any attempt to change the contents of their packets results in failure of the application. NAT does exactly that; for outgoing traffic it replaces the source private address of a hosts with an externally routable source address and replaces the corresponding private destination address for return traffic.

This change of the address on transit works with a number of applications wile others can be fixed, e.g: FTP, by using Application Layer Gateways (ALG) to also translate appropriate fields in the higher layers (e.g: TCP checksum) in order to 'hide' from the other end the fact that something has changed in the packet.

In other applications, however, the use of ALGs is either too inefficient, to be practicable (e.g: Mobile IP), or they bridge a very important part of the application in question (e.g: in IPSEC you have to trust the ALG/NAT - not always possible).

Note that the complete list of applications that break with NAT is a current NAT WG item.

This proposal provides a way to allow hosts, in networks that use NAT, to communicate with external hosts without breaking the end 2 end principle. In order to achieve the above functionality, a tunnel has to be built between the host in the private network and the NAT. The tunnel is then used to allocate an external address, out of the pool of addresses available to the NAT, as well as to route the packets outside the private network.

1.1. Assumptions

The NAT box MUST be able to handle tunnels on the interface attached to the private network. This should not be very difficult since NAT boxes are usually integrated with routers.

L2TP tunneling is assumed in this draft due to its extended functionality, but other types of tunneling MAY also be used.

It would also be helpful if the host could establish the tunnel to the NAT without human intervention. Alternatively, the tunnel MAY be statically configured and ONLY used when an application is end 2 end sensitive.

1.2. Applicability and Limitations

* This proposal does not solve problems that are inherent to NAT. In fact it does not change anything in the NAT or any other protocol. It merely bypasses NAT when the required functionality cannot be supported by NAT.

* This proposal could be used by networks that use private address space, with a small number of users that need to run applications that break through NAT. Hosts that do not need, or are not allowed by local policy, to run this kind of applications, can still use NAT in the traditional way but SHOULD NOT be allowed to use the tunnel.

* The network designer who is going to use the described mechanism needs to balance between the number of global addresses available, the total number of hosts in the private network and the number of users that are allowed to bypass the NAT at the same time.

* It is clear that when an address is allocated to a tunnel it can not be overloaded by muxing the port numbers (NAPT function)

* With this proposal NAT becomes an overloaded box. Apart from address translation, it is required to be able to handle tunnels, address allocation and potentially PPP, radius etc.

* The use of L2TP, that carries PPP packets, allows for the use of access related protocols like RADIUS, providing policy and potentially an accounting mechanism.

* NAT with the functionality described in this proposal is not transparent to the users that use the added functionality, since they need to know where to terminate the tunnel.

* NAT becomes a single point of failure for users who access it through tunnels. As far as the author understands, hot standbys may

be problematic since the tunnel configuration may be difficult to transfer.

* The use of a tunnel creates an added overhead due to tunnel headers. Header compression mechanisms for L2TP are currently investigated in [L2TPHC]

2. Requirements

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in [KEYWORDS].

3. Overview

Consider the private network in Figure 1 connected through NAT to the global Internet. Hosts A and B can communicate with Host C using NAT in order to translate their private addresses to global addresses, which are valid in the Internet.

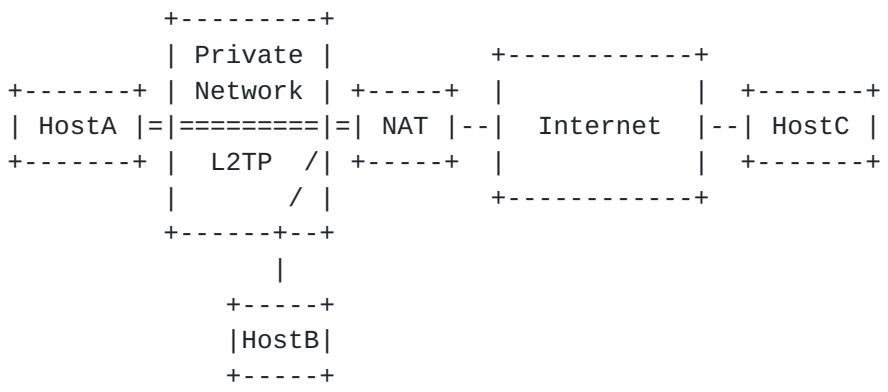

Figure 1: Tunneled connection through NAT.

Lets assume that one of the hosts, say Host A, wants to set-up an IPSEC tunnel to Host C. In order to do that Host A needs a global address that is valid end to end and is not going to be changed by the NAT box. In order to do get a global address, Host A initiates an L2TP tunnel between itself and the NAT. With normal L2TP operation, virtual interfaces are set-up in both Host A and NAT, PPP parameters are negotiated and an address, from the pool of global addresses located in the NAT, is assigned to Host A.

At the end of this procedure Host A has an IP connection running over the L2TP tunnel to the NAT, using an address valid to the global Internet. From then on, Host A can initiate a number of applications that normally would not run through the NAT, including IPSEC to Host C.

All subsequent communication to Host C is transmitted through the

L2TP tunnel in both directions and the NAT acts as a normal router without translation taking place.

The tunnel SHOULD be disconnected or at least deactivated after the session is finished and the global address MUST be returned to the NAT's pool of addresses.

4. Why Tunneling

The allocation of a globally unique address to a host in a private network is an obvious solution to networks that use NAT. This, however, creates an oxymoron in the sense that NAT is used in order avoid providing global addresses to all hosts in a network.

One could argue that if a hosts has to run applications like IPSEC frequently it might make sense to have a global address permanently allocated to it. The problem is that this is a static solution which means that even when this host does not uses its global address, the address can not be used by others. Additionally, most applications are associated with a user not a host, e.g: IPSEC is a user's decision.

It can also be argued that DHCP could be used to temporary allocate a global address. This is also problematic since the allocated address is not routable in the private domain leading to scaling problems. With Tunneling the routing problem is resolved, because the tunnel is routed on the private address.

L2TP also has the added advantage that it is configured relatively automatically and may carry PPP. The latter allows the NAT to authenticate users that want to use the added functionality applying local policy, since this is clearly an expensive function.

5. SECURITY CONSIDERATIONS

There are no security problems created by this proposal further to these described in the protocols used.

6. OPEN ISSUES

The authors do not claim to be experts on either IPSEC nor L2TP and as such, help is required to investigate and clarify the details of this proposal.

A host that wants to use the functionality described needs to know the address of the NAT. Can this be automated?

Is it possible for the tunnel to be initiated automatically when IPSEC is to be used without human intervention? Remember that the

tunnel has to be set-up and an address to be allocated before the host can initiate IPSEC.

REFERENCES

[KEYWORDS] S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.

[NAT] P. Srisuresh, et.al., The IP Network Address Translator (NAT), ftp://ietf.org/internet-drafts/draft-rfced-info-srisuresh-03.txt, September 1997

[L2TP] K. Hamzeh et.al., Layer Two Tunneling Protocol "L2TP", ftp://ietf.org/internet-drafts/draft-ietf-pppext-l2tp-08.txt, November 1997

[L2TPHC], A.J. Valencia, L2TP Header Compression (``L2TPHC''), ftp://ietf.org/internet-drafts/draft-ietf-pppext-l2tphc-01.txt, December 1997

[ROUT] C. Huitema, Routing In The Internet, 1995, Prentice Hall.

AUTHORS

George Tsirtsis
Internet Transport Group
B29 Room 129
BT Laboratoties
Martlesham Heath
IPSWICH
Suffolk IP5 3RE
England

Phone: +44 1473 640756
Fax: +44 1473 640709
e-mail: george@gideon.bt.co.uk

Alan O'Neill
Internet Transport Group
B29 Room 129
BT Laboratoties
Martlesham Heath
IPSWICH
Suffolk IP5 3RE
England

Phone: +44 1473 646459
Fax: +44 1473 640709
e-mail: alan.oneill@bt-sys.bt.co.uk

DISCLAIMER

Notice: This contribution has been prepared to assist the IETF as a basis for discussion. It is not a binding proposal on British telecommunications plc; specifically, British Telecommunications plc reserves the right to modify, delete and amend any statements contain herein.