Internet Engineering Task Force
Internet Draft
Expires in six months

R. G. Moskowitz
Chrysler Corporation
February 6, 1998

Network Address Translation issues with IPsec

<draft-moskowitz-net66-vpn-00.txt>

Status of this Memo

Abstract

This document looks at a number of issues surrounding the need for network address translation (NAT) when IPsec is used to create virtual private networks (VPN). This document only looks at simple VPNs. That is VPNs consisting of a single IPsec tunnel as compared to VPNs consisting of chained and/or nested IPsec tunnels and/or transports. It proposes a method to vastly reduce the extent that NAT is needed in a VPN.

Table of Contents

1. Introduction

This document proposes a methodology to produce simple to manage Virtual Private Networks (VPNs). It assumes that the VPNs are constructed of tunnels (IP or link layer in IP) between IP networks across an internet. If proposes a variant of Address Allocation for Private Internets [RFC 1918] to minimize the need of Network Address Translation (NAT).

The name NET66 has nothing to do with the addresses used in this document, network 7 is used as the example network that IANA would assign and manage for this VPN methodology. Rather NET66 is an allusion to the per-interstate road across the US, Route 66. Route 66 was a defined path across the western half of the US. People planned their entrance and exit points from Route 7. Likewise, NET66 means some planning for entrance and exit from the VPN, but the expectation of a smooth trip over it.

1.1 Requirements

The Design objectives are:

All servers in the VPN are globally addressed.

The servers need not be globally routable:

Server addresses need to be routable in the server s network.
Server addresses need to be default or statically routed to the VPN gateway on the client s network.

The clients need a name they can resolve that yields an address that will route to the client s gateway.

The clients gateway needs to discover the server s gateway s address with only the server s address as a clue.

2. Network Description

The following diagram will be used in this discussion as a sample VPN of a community of interest.

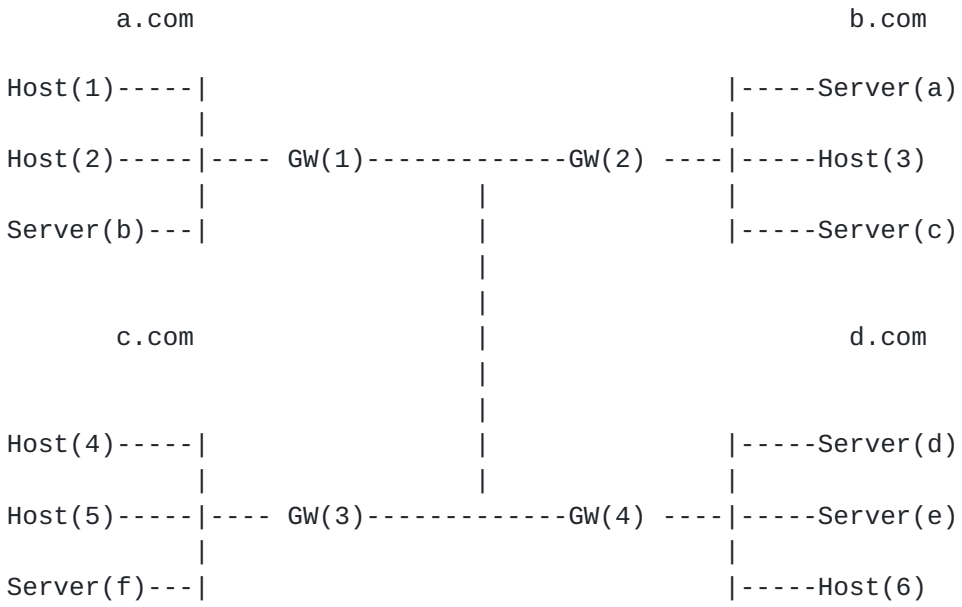

3. Addressing for VPN Servers

If Srva.b.com has a global address then Host1.a.com can resolve the name to the address. A.com will have to have a route for this address to GW1. If all the VPN accessible servers are addressed out of the same block, then a.com only needs one route to GW1 for all servers accessed by host1.a.com. Since all of the traffic to the servers is within the VPN tunnels these addresses need not be routable

on the internet, they only need be registered in the internet s DNS.

3.1 Sample DNS entries for VPN Servers

Assume that the block of addresses for VPN Servers is 7.0.0.0/8. If the 4 companies in section 2 received an allocation of 2 addresses the DNS entries might look like:

```
Srvb.a.com          IN   A    7.0.1.1
Srva.b.com          IN   A    7.0.2.1
Srvc.b.com          IN   A    7.0.2.2
Srvf.c.com          IN   A    7.0.3.1
Srvd.d.com          IN   A    7.0.4.1
```

3.2 Internal addressing and routing in VPN Server networks

The addresses in 3.1 would be used both in the public DNS and as the actual address on the servers. The router adjacent to the server would advertise a unique route for the server (e.g. 7.0.1.1/32). This way the packets will reach the server as they emerge from the gateway without any address translation required.

3.3 Internal routing in VPN Client networks

The client networks (e.g. a.com for host1.a.com) advertise a route of 7.0.0.0/8 directed to their gateway.

3.4 Sample Process flow

Host1.a.com needs to connect to a TCP application on Srva.b.com. Host1 performs a DNS lookup and gets the address 7.0.2.1. Host1 (assume its address is 192.168.50.2) sends out a TCP SYN sourced from 192.168.50.2, destined for 7.0.2.1.

A.com s network routes this packet into GW1 since it has a route to 7.0.0.0/8. GW1 (whose public address is 199.175.30.1) has a policy (see sec 4) that a destination of 7.0.2.1 is tunneled to 208.50.1.1 (GW2 s public address). GW1 places the packet it received into an IP packet sourced from 199.175.30.1, destined for 208.50.1.1.

The internet routes this packet to GW2. GW2 removes the internal packet. Since the source address is a private

address, GW2 MUST translate this address to an address usable within b.com (for example for the pool of 172.17.1.0/24). This requirement is based on the likelihood of GW2 having two or more tunnels from different hosts, all having the same private address. If the source address had been a public address, GW2 COULD have left the address unchanged if b.com s network is able to route that address (as a destination address) back to GW2 (this is a site option). GW2 now delivers the packet into b.com s network.

B.com s network routes this packet to Srva since there is a route to 7.0.2.1. Srva reverses the source and destination addresses and creates an SYN/ACK. B.com s network routes this packet back to GW2. GW2 reverses the address translation on the destination address (if it performed one) and places the packet into an IP packet sourced from 208.50.1.1, destined for 199.175.30.1 (that is GW1).

The internet delivers this packet to GW1. It removes the inner packet and puts it on a.com s network. A.com s network delivers the packet to Host1.

4. Tunnel discovery the KX record

A key point in section 3.4 is GW1 knowning that a packet destined for 7.0.2.1 was to be tunneled to 208.50.1.1. This could be done via manual configuration; that is GW1 s administrator had set up a rule accordingly. This method is easy to deploy on a small scale, but is not usable in a large context. An alternative is for GW1 to be able to process DNS KX records [RFC 2230]. Consider the following DNS records:

```
Srva.b.com.                IN   A    7.0.2.1
                           IN   KX   10, gw2.b.com.
1.2.0.7.in-addr.arpa.      IN   A    Srva.b.com.
gw2.b.com.                 IN   A    208.50.1.1
```

4.1 KX Process flow

A configuration option on GW1 starts the KX processing if the destination address is 7.0.0.0/8. GW1 does a reverse lookup on the destination address of 7.0.2.1 and gets Srva.b.com. It next does a lookup for a KX record for Srva.b.com and gets gw2.b.com. Finally it does a lookup up on gw2.b.com and gets 208.50.1.1 as the tunnel end point.

Since KX processing s purpose is to retrieve VPN configuration information, all of these queries MUST be secured with DNSSEC. This, plus the three DNS queries might be considered as a drawback to using the KX record. Another drawback in the KX record is that it does not supply any tunnel policy information, like a hint at what type of tunneling technology or parameters to use.

4.2 Alternatives to the KX record

There are many alternatives to the KX record. SRVLOC, LDAP directory, and Attribute Certificates have all been proposed. Attribute certificates have the security needs stated above built into them. A full analysis is needed to select a better alternative to the KX approach.

5. IANA functions

IANA will reserve a large block of addresses, a /8 is recommended (and might not be enough until IPv6 is deployed eliminating the need of NET66). IANA will develop the processes and procedures for an organization to license addresses from this block for their VPN accessible servers.

6. Security Considerations

Network address translation, in conjunction with IPsec makes some large assumptions of trust. Intermediate systems are changing IP addresses on behalf of other systems.

The KX record relies on the deployment of DNSSEC, without this it cannot be trusted. In fact the whole global server address block and its reverse lookup block needs to be secured with DNSSEC as well.

7. References

[RFC 1918] Y. Rekhter, B. Moskowitz, D. Karrenberg, G. J. de Groot & E. Lear Address Allocation for Private Internets
[RFC 2230] R. Atkinson Key Exchange Delegation Record for the DNS

8. Acknowledgments

This document is based on discussions with Mark Johnson and Mike Papais of Chrysler Corporation, along with a host of others at the the Automotive Industry Action Group (AIAG) and within the IETF community.

9. Author's Addresses

Robert G. Moskowitz
rgm@icsa.net
ICSA, Inc.