Network Working Group Michael Richardson mcr@sandelman.ottawa.on.ca
INTERNET-DRAFT SSH Communications Security
draft-richardson-ipsec-icmp-filter.txt v1.0, 16 July 1997
Expires in six months

Why traceroute can not work through IPsec gateways

Status of This memo

Abstract

This document describes the problem of doing diagnostics through IPsec gateways (VPNs). If the gateways implement their policies to the letter, then diagnostics are not possible.

Table of Contents

1. Introduction

1.1. Definition of terminology

Here is a network of two security gateways, a client node and a server node.

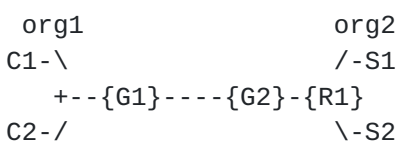

C1 is a host.
C2 is another host.
G1/G1 are security gateways.
S1 is a host.
S2 is a host.
R1 is a router.

There are per-host SA's linking C1<->S1, C2<->S2, and also C1<->S2.

One does a traceroute from C1 to S.

One expects to see:

```
traceroute to S1 (192.168.32.71), 30 hops max, 40 byte packets
 1  G1  2.323 ms  2.323 ms  2.323 ms
 2  G2  3.456 ms  3.456 ms  3.456 ms
 3  R1  8.456 ms  8.456 ms  8.456 ms
 4  S1  5.678 ms  5.678 ms  5.678 ms
```

The first return is from G1. The second return is from G2, the final one is from S. Let's examine the details of the ICMP datagrams that are received by C1.

1. A datagram from G1 to C1, traverses a protected (unencrypted network). No problem, there is no SA to restrict traffic between G1

and C1.

2. A datagram from G2 to C1. This is a problem: The SA between G1 and G2

is for datagrams between C1 and S1, but the datagram doesn't fit into that pattern. It is reasonable to assume that the filter code could be persuaded to accept this packet. It is from a node that is trusted (G2), which could easily send a spoofed datagram (claiming to be from S1) through tunnel if it wanted to.

The question remains: which SA should be used? Clearly, it should be the one linking C1 and S1, not the one linking C1 and S2. The two SA's could have very different privacy attributes. One must pick the right SA bundle. This problem is solvable at this level.

3. **A datagram from R1 to C1. This is an even worse problem. The gateway has little ability to know if R1 is even legitimately a router on** which datagrams to S1 must travel. Further, there is now no record in the outer ip header as to which SA to use.

4. **A datagram from S2 to G2. No problem, the SA covers this already.**

2. Options HEADER-2

2.1. No ICMP?

One possible solution is to give up on moving ICMP datagrams at all.

2.2. ICMP from preset list

Maybe it is enough to accept ICMP datagrams from a preconfigured list of routers. This list would include the IPsec gateway (G2), and the list would have to be passed to G1.

2.3. **The header inside the ICMP packet gives the truth. This could be** used to determine if the packet had appropriate source/destinations. Examine the ICMP HEADER-2

2.4. ICMP soft state?

The ICMP datagram carries some 28 bytes (plus options) of the original datagram. The destination/source/id field ought to be unique. The gateways could take arrival of an ICMP datagram with some destination as an indication to start recording datagram ids (i.e. accumulating soft state). A retransmission occurs, and then the ICMP can be matched to an actual IP datagram.

This is not very secure, as it can be spoofed by any node on org2's network. The assumed requirement for host based SA's, implies a distrust of org2's network.

3. References:

RFC-1825

R. Atkinson, "Security Architecture for the Internet Protocol", RFC-1825, August 1995.

RFC-1191
J. Mogul, S. Deering, "Path MTU Discovery", RFC-1191, November 1990.

RFC-1812
F. Baker, "Reqirements for IP Version 4 Routers", RFC-1812, June 1995

3.1. Author's Address

Michael C. Richardson
Sandelman Software Works Corp.
152 Rochester Street
Ottawa, ON K1R 7M4
Canada

Telephone: +1 613 233-6809
EMail: mcr@sandelman.ottawa.on.ca

3.2. Expiration and File Name

This draft expires January 9, 1997

Its file name is draft-richardson-ipsec-icmp-filter-00.txt