

IPsec Working Group
INTERNET-DRAFT
Category: Informational
<[draft-ietf-ipsec-nat-reqts-00.txt](#)>
[18](#) June 2001

Bernard Aboba
Microsoft

IPsec-NAT Compatibility Requirements

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at
<http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at
<http://www.ietf.org/shadow.html>.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

Perhaps the most common use of IPsec is in providing virtual private networking capabilities. One very popular use of VPNs is to provide tele-commuter access to the corporate Intranet. Today NATs are widely deployed in home gateways, as well as in other locations likely to be used by tele-commuters, such as hotels. The result is that IPsec-NAT incompatibilities have become a major barrier to deployment of IPsec in one of its principal uses. This draft describes known incompatibilities between NAT and IPsec, and describes the requirements for addressing them.

[1](#). Introduction

Perhaps the most common use of IPsec [[6](#)] is in providing virtual private networking capabilities. One very popular use of VPNs is to provide

INTERNET-DRAFT

IPsec-NAT Compatibility Reqts.

18 June 2001

tele-commuter access to the corporate Intranet. Today NATs [8]-[9] are widely deployed in home gateways, as well as in other locations likely to be used by tele-commuters, such as hotels. The result is that IPsec-NAT incompatibilities have become a major barrier to deployment of IPsec in one of its principal uses. This draft describes known incompatibilities between NAT and IPsec, and describes the requirements for addressing them.

[1.1.](#) Requirements language

In this document, the key words "MAY", "MUST", "MUST NOT", "optional", "recommended", "SHOULD", and "SHOULD NOT", are to be interpreted as described in [2].

Please note that the requirements specified in this document are to be used in evaluating protocol submissions. As such, the requirements language refers to capabilities of these protocols; the protocol documents will specify whether these features are required, recommended, or optional. For example, requiring that a protocol support confidentiality is NOT the same thing as requiring that all protocol traffic be encrypted.

A protocol submission is not compliant if it fails to satisfy one or more of the MUST or MUST NOT requirements for the capabilities that it implements. A protocol submission that satisfies all the MUST, MUST NOT, SHOULD and SHOULD NOT requirements for its capabilities is said to be "unconditionally compliant"; one that satisfies all the MUST and MUST NOT requirements but not all the SHOULD or SHOULD NOT requirements for its protocols is said to be "conditionally compliant."

[2.](#) Known incompatibilities between NAT and IPsec

The incompatibilities between NAT and IPsec are numerous, ranging from the obvious to the subtle. Some of the known incompatibilities include:

- a) Incompatibility between IPsec AH [3] and NAT. Since the AH header incorporates the IP source and destination addresses in the keyed message integrity check, NAT or reverse NAT devices making changes to address fields will invalidate the message integrity check. Since IPsec ESP [4] does not incorporate the IP source and destination addresses in its keyed message integrity check, this issue does not arise for ESP.

- b) Incompatibility between checksums and NAT. TCP/UDP/SCTP checksums have a dependency on the IP source and destination addresses through inclusion of the "pseudo-header" in the calculation. As a result, where checksums are calculated and checked on receipt, they will be invalidated by passage through

a NAT or reverse NAT device.

As a result, IPsec ESP will only pass unimpeded through a NAT if TCP/UDP/SCTP protocols are not involved (as in IPsec tunnel mode or IPsec/GRE), or checksums are not calculated (as is possible with IPv4 UDP). As described in [13], TCP checksum calculation and verification is required in IPv4. UDP/TCP checksum calculation and verification is required in IPv6.

Note that since transport mode IPsec traffic is integrity protected and authenticated using strong cryptography, modifications to the packet can be detected prior to checking UDP/TCP/SCTP checksums. Thus, checksum verification only provides assurance against errors made in internal processing.

- c) Incompatibility between IKE address identifiers and NAT. Where IP addresses are used as identifiers in IKE MM [7] or QM, modification of the IP source or destination addresses by NATs or reverse NATs will result in a mismatch between the identifiers and the addresses in the IP header. As described in [7], IKE implementations are required to discard such packets.

In order to avoid use of IP addresses as IKE MM and QM identifiers, userIDs and FQDNs can be used instead. Where user authentication is desired, an ID type of ID_USER_FQDN can be used, as described in [5]. Where machine authentication is desired, an ID type of ID_FQDN can be used. In either case it is necessary to verify that the proposed identity matches that enclosed in the certificate. However, while use of USER_FQDN or FQDN identity types is possible within IKE, there are usage scenarios (e.g. SPD entries describing subnets) that cannot be accommodated this way.

- d) Incompatibility between fixed IKE destination ports and NAPT. Where multiple hosts behind the NAPT initiate IKE SAs to the same responder, a mechanism is needed

to allow the NAT to demultiplex the incoming IKE packets. This is typically accomplished by translating the IKE UDP source port. While it is permissible to float the IKE UDP source port, this can result in unpredictable behavior during re-keys. Unless the floated source port is used as the destination port for the re key, the NAT may not be able to send the re-key packets to the correct destination.

- e) Incompatibilities between IKE cookie usage and NAT.
Today some NAT implementations attempt to use IKE cookies to de-multiplex incoming IKE traffic. As with source-port

Aboba

Informational

[Page 3]

INTERNET-DRAFT

IPsec-NAT Compatibility Reqts.

18 June 2001

de-multiplexing, IKE cookie de-multiplexing also results in problems with re-keying, since re-keys typically will not use the same cookies as the earlier traffic.

- f) Incompatibilities between overlapping SPD entries and NAT.
Where hosts behind a NAT negotiate overlapping SPD entries with the same destination in IKE QM, packets may be sent down the wrong IPsec SA. This occurs because to the sender, the IPsec SAs appear to be equivalent, since they exist between the same endpoints and can be used to pass the same traffic.
- g) Incompatibilities between IPsec SPI selection and NAT.
Since IPsec ESP traffic is encrypted and thus opaque to the NAT, the NAT must use elements of the IP and IPsec header to demultiplex incoming IPsec traffic. The combination of the destination IP address, security protocol (AH/ESP) and IPsec SPI is typically used for this purpose.

However, since the outgoing and incoming SPIs are chosen independently, there is no way for the NAT to determine what incoming SPI corresponds to what destination host merely by inspecting outgoing traffic. Thus, were two hosts behind the NAT to attempt to bring up IPsec SAs to the same destination simultaneously, it is possible that the NAT will send the incoming IPsec packets to the wrong destination.

Note that this is not an incompatibility with IPsec per

se, but rather with the way it is typically implemented. With both AH and ESP, the receiving host specifies the SPI to use for a given SA. At present, the combination of Destination IP, SPI, and Security Protocol (AH, ESP) uniquely identifies a Security Association. This means that the receiving host can select SPIs such that it has one Security Association (SA) with (SPI=470, Dest IP=1.2.3.4) and a different Security Association with (SPI=470, Dest IP=2.3.4.5).

It is also possible for the receiving host to allocate a unique SPI to each unicast Security Association. In this case, the Destination IP Address need only be checked to see if it is "any valid unicast IP for this host", not checked to see if it is the specific Destination IP address used by the sending host. This approach is completely backwards compatible and only requires the particular receiving host to make a change to its SPI allocation and IPsec_esp_input() code.

- h) Incompatibilities between embedded IP addresses and NAT. Since the payload is integrity protected, any IP addresses enclosed within IPsec packets will not be translatable by a NAT. Protocols that utilize embedded IP addresses include FTP, IRC, SNMP, LDAP, H.323, SIP and many games.

[2.1](#). Why IPsec tunnel mode works

Given the incompatibilities described above, it might seem unlikely that any IPsec/IKE conversation could survive passage through a NAT. However, IPsec tunnel mode clients [\[15\]](#) are capable of traversing NATs under limited conditions:

- [1] IPsec ESP. IPsec ESP tunnels do not cover the outer IP header within the authentication hash, and so will not suffer hash invalidation due to address translation. IPsec tunnels also need not be concerned about checksum invalidation (unlike L2TP).
- [2] Address validation. Most current IPSEC tunnel mode implementations do not perform source address validation so that incompatibilities between IKE identifiers and source addresses will not be detected. This introduces security vulnerabilities as described in the

security considerations section.

- [3] SPD entries. Most IPsec tunnel mode clients negotiate "any to any" SPDs, which are not invalidated by address translation. This effectively precludes use of SPDs for filtering of allowed tunnel traffic.
- [4] Single client operation. With only a single client behind a NAT, there is no risk of overlapping SPDs. Since the NAT will not need to arbitrate between competing clients, there is also no risk of re-key mis-translation, or improper incoming SPI or cookie de-multiplexing.

3. Requirements for IPsec-NAT compatibility

The goal of an IPsec-NAT compatibility solution is to expand the range of usable IPsec functionality beyond that available in an NAT-compatible IPsec tunnel mode solution described above.

In evaluating a solution to IPsec-NAT incompatibility, the following criteria should be kept in mind:

Deployability

Since IPv6 will address the address scarcity issues that frequently lead to use of NATs with IPv4, the IPsec-NAT compability issue is a transitional problem that needs to be

solved in the timeframe prior to widespread deployment of IPv6. Therefore, to be useful an IPsec-NAT compatibility solution MUST be deployable on a shorter time scale than IPv6.

Since IPv6 deployment requires changes to routers as well as hosts, a IPsec-NAT compatibility solution which requires changes to both routers and hosts will be deployable on approximately the same time scale as IPv6. Thus, an IPsec-NAT compatibility solution SHOULD require changes only to hosts, and not to routers.

Among other things, this implies that communication between the host and the NAT MUST NOT be required by an IPsec-NAT compatibility solution, since existing NATs cannot meet such a requirement.

Telecommuter scenario

Since a typical telecommuter is only interested in obtaining access to the corporate Intranet for itself, an IPsec-NAT compatibility solution need not enable gateway-gateway connectivity. Thus, it can be assumed that negotiated SPD entries will only refer to the communication endpoints, and there is no need to support negotiation of SPD entries involving subnets. to

Scaling An IPsec-NAT compatibility solution should be capable of being deployed within an installation consisting of thousands of telecommuters. In this situation, it is not possible to assume that only a single host is communicating with a given destination at a time. Thus, an IPsec-NAT compatibility solution **MUST** address the issue of overlapping SPD entries and de-multiplexing of incoming packets.

Mode support

At a minimum, an IPsec-NAT compatibility solution **MUST** support passage of IPsec ESP tunnel mode through a NAT. Since IPsec transport mode is used for tunneling protocols such as L2TP [1], an IPsec-NAT compatibility solution **SHOULD** support IPsec transport mode using ESP, at least for protection of UDP traffic where no embedded IP addresses are present. Since passage of AH through a NAT is not possible in any mode, there is no need for an IPsec-NAT compatibility solution to attempt to address this.

Interoperability

An IPsec-NAT compatibility solution **MUST** be interoperable with existing IPsec implementations. Thus, existing IPsec implementations **MUST** be able to communicate with an IPsec-NAT

compatible implementation in the case where no NAT is present. This implies that an IPsec-NAT compatibility solution **MUST** be backwards-compatible with IPsec as defined in [3]-[7], and **SHOULD** be able to detect the presence of a NAT so that the required changes for NAT compatibility will only be used when necessary.

For example, it may be possible to enable ISAKMP to pass

information about each host's perception of its own IP address, in order to make key management aware of the presence of the NAT and facilitate the use of standard key management methods through a NAT to support ESP/AH.

Security An IPsec-NAT compatibility solution MUST NOT introduce additional security vulnerabilities into IKE. For example, an acceptable solution must demonstrate that it introduces no new denial of service or spoofing vulnerabilities.

[4.](#) Existing solutions

[4.1.](#) RSIP

RSIP, described in [[10](#)]-[[11](#)], includes mechanisms for IPsec traversal, as described in [[12](#)]. By enabling host-gateway communication, RSIP addresses issues of IPsec SPI de-multiplexing as well as SPD overlap. It is thus suitable for use in enterprise as well as home networking scenarios. By enabling hosts behind a NAT to share the external IP address of the gateway, this approach is compatible with protocols including embedded IP addresses.

By tunneling IKE and IPsec packets, RSIP avoids changes to the IKE and IPsec protocols, although major changes are required to host IKE and IPsec implementations to retrofit them for RSIP-compatibility. It is thus compatible with all existing protocols (AH/ESP) and modes (transport and tunnel).

In order to handle de-multiplexing of IKE re-keys, RSIP requires floating of the IKE source port, as well as re-keying to the floated port. As a result, inter-operability with existing IPsec implementations is not assured.

RSIP does not satisfy the deployability requirements for a IPsec-NAT compatibility solution because an RSIP-enabled host requires a corresponding RSIP-enabled gateway in order to establish an IPsec SA with another host. Since RSIP requires changes only to clients and routers and not to servers, it is less difficult to deploy than IPv6. However, for vendors, implementation of RSIP requires a substantial

fraction of the resources required for IPv6 support. Thus, RSIP solves a

"transitional" problem on a long-term time scale, which is not useful.

5. Security considerations

By definition, IPsec-NAT compatibility requires that hosts and routers implementing IPsec be capable of securely processing packets whose IP headers are not cryptographically protected. A number of issues arise from this that are worth discussing.

Since IPsec AH cannot pass through a NAT, one of the side effects of providing an IPsec-NAT compatibility solution may be for IPsec ESP with null encryption to be used in place of AH where a NAT exists between the source and destination. However, it should be noted that ESP with null encryption does not provide the same security properties as AH. For example, there are security risks relating to IP source routing that are precluded by AH, but not by ESP with null encryption.

In addition, since ESP with any transform does not protect against source address spoofing, some sort of source IP address sanity checking needs to be performed. The importance of the anti-spoofing check is not widely understood. There is normally an anti-spoofing check on the Source IP Address as part of IPsec_{esp,ah}_input(). This ensures that the packet originates from the same address as that was claimed within the original IKE MM and QM security associations. When a receiving host is behind a NAT, this check might not strictly be meaningful for unicast sessions, whereas in the Global Internet this check is important for tunnel-mode unicast sessions to prevent a spoofing attack described in [\[14\]](#).

Let us consider two hosts, A and C, both behind (different) NATs, who negotiate IPsec tunnel mode SAs to router B. Hosts A and C may have different privileges; for example, host A might belong to an employee trusted to access much of the corporate Intranet, while C might be a contractor only authorized to access a specific web site.

If host C sends a tunnel mode packet spoofing A's IP address, as the source, it is important that this packet not be accorded the privileges corresponding to A. If authentication and integrity checking is performed, but no anti-spoofing check (verifying that the originating IP address corresponds to the SPI) then host C may be allowed to reach parts of the network that are off-limits. As a result, an IPsec-NAT compatibility scheme MUST provide some degree of anti-spoofing protection.

6. Acknowledgments

Thanks to Steve Bellovin of AT&T Research, William Dixon of Microsoft, Ran Atkinson of Extreme Networks and Daniel Senie for useful discussions of this problem space.

7. References

- [1] Townsley, W., Valencia, A., Rubens, A., Pall, G., Zorn, G., and Palter, B., "Layer Two Tunneling Protocol L2TP", [RFC 2661](#), August 1999.
- [2] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [3] Kent, S., Atkinson, R., "IP Authentication Header", [RFC 2402](#), November 1998.
- [4] Kent, S., Atkinson, R., "IP Encapsulating Security Payload (ESP)", [RFC 2406](#), November 1998.
- [5] Piper, D., "The Internet IP Security Domain of Interpretation of ISAKMP", [RFC 2407](#), November 1998.
- [6] Atkinson, R., Kent, S., "Security Architecture for the Internet Protocol", [RFC 2401](#), November 1998.
- [7] Harkins, D., Carrel, D., "The Internet Key Exchange (IKE)", [RFC 2409](#), November 1998.
- [8] Srisuresh, P., and Egevang, K., "Traditional IP Network Address Translator (Traditional NAT)", [RFC 3022](#), January 2001.
- [9] Srisuresh, P. and Holdredge, M., "IP Network Address Translator (NAT) Terminology and Considerations", [RFC 2663](#), August 1999.
- [10] Borella, M., Lo, J., Grabelsky, D., Montenegro, G., "Realm Specific IP: A Framework", Internet draft (work in progress), [draft-ietf-nat-rsip-framework-05.txt](#), July 2000.
- [11] Borella, M., Grabelsky, D., Lo, J., Taniguchi, K., "Realm Specific IP: Protocol Specification", Internet draft (work in progress), [draft-ietf-nat-rsip-protocol-07.txt](#), July 2000.
- [12] Montenegro, G., Borella, M., "RSIP Support for End-to-End IPsec", Internet draft (work in progress), [draft-ietf-nat-rsip-](#)

INTERNET-DRAFT

IPsec-NAT Compatibility Reqs.

18 June 2001

- [13] Information Sciences Institute, "Transmission Control Protocol", [RFC 793](#), September 1981.
- [14] Kent, S., "Authenticated Source Addresses", IPsec WG Archive (<ftp://ftp.ans.net/pub/archive/IPsec>), Message-Id: <v02130517ad121773c8ed@[128.89.0.110]>, January 5, 1996.
- [15] Patel, B., Aboba, B., Kelly, S., Gupta, V., "DHCPv4 Configuration of IPsec Tunnel Mode", Internet draft (work in progress), [draft-ietf-ipsec-dhcp-12.txt](#), May 2001.

[8.](#) Authors' Addresses

Bernard Aboba
Microsoft Corporation
One Microsoft Way
Redmond, WA 98052

Phone: +1 425 936 6605
EMail: bernarda@microsoft.com

[9.](#) Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Information on the IETF's procedures with respect to rights in standards-track and standards-related documentation can be found in [BCP-11](#). Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights

which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

10. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to

Aboba

Informational

[Page 10]

INTERNET-DRAFT

IPsec-NAT Compatibility Reqts.

18 June 2001

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."

11. Expiration Date

This memo is filed as <[draft-ietf-ipsec-nat-reqts-00.txt](#)>, and expires January 10, 2002.

