

ipsecme
Internet-Draft
Intended status: Informational
Expires: September 10, 2015

A. Antony
Phenome Networks
J. Gilmore

P. Wouters
Red Hat
March 09, 2015

NAT-Traversal support for Opportunistic IPsec
draft-antony-ipsecme-oppo-nat-00

Abstract

This document specifies how to support NATed IPsec peers for use with Opportunistic IPsec.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on September 10, 2015.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
1.1.	Conventions Used in This Document	2
2.	The pre-NAT IP address conflict	2
3.	Creating the NAT workaround on the initiator	3
3.1.	The Initial Exchange	3
3.2.	Handling the NAT problem on the initiator	4
3.3.	Handling the NAT problem on the responder	4
3.4.	Implementation details	4
4.	Creating the NAT workaround on the responder	5
5.	Security Considerations	5
6.	Acknowledgments	5
7.	IANA Considerations	5
8.	References	5
8.1.	Normative References	5
8.2.	Informative References	5
	Authors' Addresses	6

[1.](#) Introduction

The Internet Key Exchange Protocol version 2 (IKEv2), specified in [\[RFC7296\]](#), provides a way to negotiate tunnel mode IPsec SA's for peers behind NAT. It does not provide guidance on how to resolve the problem of multiple peers using the same pre-NAT IP address.

Responder assigned IP addresses for NATed peers also do not fully resolve the address conflict when the NATed peers are deploying opportunistic IPsec to many remote endpoints. Additional complexity of configuring many source IP addresses on the NATed peers is undesirable.

This problem is expected to be a significant issue for large scale Opportunistic IPsec deployments.

The goal of this draft is to put the burden of the additional effort required for NAT onto the host that is NAT'ed. In a hypothetical future with no NAT, no residual protocol changes would remain.

[1.1.](#) Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this

document are to be interpreted as described in [[RFC2119](#)].

[2.](#) The pre-NAT IP address conflict

Antony, et al.

Expires September 10, 2015

[Page 2]

Internet-Draft

NAT-T for Opportunistic IPsec

March 2015

A peer that is behind NAT is only aware of its own native IP address. While the IKEv2 NATD payloads allow the peer to find out it is behind NAT, it either has to use its own native IP address or a remote peer assigned IP address to build IPsec SA's.

Two initiators behind different NAT routers can have the same pre-NAT IP address. When these initiators attempt to setup an IPsec SA to the same responder, there will be a conflict on the responder. When the first initiator connects, the responder will install an IPsec SA for for the specific pre-NAT IP. For the second initiator, the responder cannot install an IPsec SA with the same source and destination selectors without creating a conflict. Even if it could install two identical IPsec SAs, it still needs a mechanism to decide between these two IPsec SA's. While connection tracking might be used to distinguish reply packets, there is no method for initiating a connection from the responder to one of the initiators without somehow specifying for which of the initiators the packet is meant.

When initiators behind NAT let the remote peer assign their IP to use for building the IPsec SA, the problem reverses. Initiators behind NAT can setup many IPsec SA's to different remote peers, and those remote peers might use the conflicting IP addresses. Additionally, using many IP addresses requires that the host or applications need to be aware of which source IP to use for which remote peer.

[3.](#) Creating the NAT workaround on the initiator

In the below workflow, the following example IP addresses are used:

The initiator ("road") has a private IP address of 192.1.3.209

The NAT router has an internal (private) IP address of 192.1.3.254

The NAT router has an external (public) IP address of 192.1.2.254

The responder ("east") has a public IP address of 192.1.2.42

The responder ("east") has a private IP address pool of 10.1.2.0/24

Tunnel mode is used for all IPsec SA's.

[3.1.](#) The Initial Exchange

In IKEv2, the IKE_INIT exchange allows the initiator and responder to detect the presense of a NAT. In this case both "road" and "east" become aware of the NAT in front of "road".

Antony, et al.

Expires September 10, 2015

[Page 3]

Internet-Draft

NAT-T for Opportunistic IPsec

March 2015

In IKEv2, the initiator can signal via the TSi payload that it is willing to receive a responder-assigned IP address for use in the IPsec SA. The responder needs to be configured with an addresspool from which it assigns unique IP addresses to the connecting initiators.

"road" sends a TSi(0.0.0.0/0) and TSr(192.1.2.42/32) request to "east" in the IKE_AUTH Exchange.

"east" assigns the first free IP address from its pool - 10.1.2.1 - to this prospective client and sends a TSi(10.1.2.1/32) and TSr(192.1.2.42/32) back to "road".

The IKE_AUTH Exchange completes as normal. This could be an authenticated exchange or an exchange without authentication using AUTH_NULL as specified in [[IKE-AUTH-NULL](#)].

[3.2.](#) Handling the NAT problem on the initiator

"road", upon receiving east's TSi/TSr - and a successfull completion of the IKE_AUTH Exchange - two tasks:

1. It installs the negotiated inbound and outbound IPsec SAs (10.1.2.1/32 <=> 192.1.2.42/32)
2. It installs a destination based NAT rule for 192.1.3.209 -> 192.1.2.42/32 ==> 10.1.2.1/32 -> 192.1.2.42/32

It ensures that the destination based NAT rule is processed before

IPsec processing begins.

[3.3.](#) Handling the NAT problem on the responder

The responder has no specific handling it needs to perform apart from a willingness to assign IP addresses if an initiator requests one via the special TSi(0.0.0.0/0). If the initiator sends a TSi() that matches the IP address of the IKE message, the responder should agree to use that address instead of assigning one from its addresspool.

[3.4.](#) Implementation details

The NAT rule can be implemented in the Networking subsystem of the host kernel, but this specific IPsec-NAT rule could also be implemented as part of the IPsec subsystem.

Antony, et al.

Expires September 10, 2015

[Page 4]

Internet-Draft

NAT-T for Opportunistic IPsec

March 2015

The initiator host and any application running on the initiator should not need to be aware of this construct. The responder assigned IP address does not need to be configured on the host. It only needs to exist in the NAT rule and the IPsec SA.

[4.](#) Creating the NAT workaround on the responder

[placeholder]

[5.](#) Security Considerations

This NAT mapping method MUST only be used for host-to-host IPsec tunnels. It MUST NOT be used for net-to-net IPsec tunnels.

An address from the addresspool should not be re-used quickly to avoid sending traffic meant for one initiator to another initiator.

[6.](#) Acknowledgments

None so far

[7.](#) IANA Considerations

This Internet Draft includes no request to IANA.

[8.](#) References

[8.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), December 2005.
- [RFC7296] Kaufman, C., Hoffman, P., Nir, Y., Eronen, P., and T. Kivinen, "Internet Key Exchange Protocol Version 2 (IKEv2)", STD 79, [RFC 7296](#), October 2014.
- [IKE-AUTH-NULL]
Smyslov, Valery. and Paul. Wouters, "The NULL Authentication Method in IKEv2 Protocol", [draft-ietf-ipsecme-ikev2-null-auth](#) (work in progress), February 2015.

[8.2.](#) Informative References

- [RFC7258] Farrell, S. and H. Tschofenig, "Pervasive Monitoring Is an Attack", [BCP 188](#), [RFC 7258](#), May 2014.

Antony, et al. Expires September 10, 2015 [Page 5]

Internet-Draft NAT-T for Opportunistic IPsec March 2015

- [RFC7435] Dukhovni, V., "Opportunistic Security: Some Protection Most of the Time", [RFC 7435](#), December 2014.

Authors' Addresses

Antony Antony
Phenome Networks

Email: antony@phenome.org

John Gilmore

Email: gnu@toad.com

Paul Wouters
Red Hat

Email: pwouters@redhat.com