Network Working Group	X. Jiang
Internet-Draft	Huawei
Intended status: Informational	P. Matthews
Expires: May 13, 2008	Avaya
	S. Dawkins, Ed.
	Huawei (USA)
	November 10, 2007

TOC

Proposed Host Identity Protocol (HIP) Checksum Coverage
draft-dawkins-hip-checksum-coverage-00

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.
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/1id-abstracts.txt.
The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.
This Internet-Draft will expire on May 13, 2008.

Abstract

This specification suggests two changes to the Host Identity Protocol (HIP) checksum calculation. Specifically, only the HIP header and payload fields would be included in the checksum calculation, and the CRC32c algorithm would be used to compute the checksum. The HIP version number would be incremented if these suggestions are accepted, to reflect a different checksum algorithm.

Table of Contents

---

1. Introduction

TOC

This specification suggests two changes to the Host Identity Protocol (HIP) checksum calculation. Each of these changes should be considered separately:

1. A change to the checksum coverage in HIP, and

2. Replacing the 16-bit checksum with a CRC32c checksum.

The HIP version number would be incremented if these suggestions are accepted, to reflect a different checksum algorithm.
The checksum calculation described in this specification is used only when HIP is carried directly over IP. Considerations for checksum calculations when an intervening transport protocol is used are left to [I-D.ietf-hip-nat-traversal] (Komu, M., Henderson, T., Tschofenig, H., Melen, J., and A. Keranen, “Basic HIP Extensions for Traversal of Network Address Translators,” October 2009.).

---

2. Terminology used in this document

TOC

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] (Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” March 1997.).

3. Background for this Proposal

TOC

In section 5.1.1, the HIP base protocol draft [I-D.ietf-hip-base] (Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson, “Host Identity Protocol,” October 2007.) now specifies the following:

> Since the checksum covers the source and destination addresses in the IP header, it must be recomputed on HIP-aware NAT devices.
>
> If IPv6 is used to carry the HIP packet, the pseudo-header [RFC2460] contains the source and destination IPv6 addresses, HIP packet length in the pseudo-header length field, a zero field, and the HIP protocol number (see Section 4) in the Next Header field. The length field is in bytes and can be calculated from the HIP header length field: (HIP Header Length + 1) * 8.
>
> In case of using IPv4, the IPv4 UDP pseudo header format [RFC0768] is used. In the pseudo header, the source and destination addresses are those used in the IP header, the zero field is obviously zero, the protocol is the HIP protocol number (see Section 4), and the length is calculated as in the IPv6 case.

Although the use of pseudo-headers in transport checksums has been universal practice for decades, this design choice is proving problematic for several applications that perform HIP-level packet relays. For example:

- There are two active proposals in the HIP working group itself that require forwarding of HIP packets, "Host Identity Protocol (HIP) Rendezvous Extension" ([I-D.ietf-hip-rvs] (Laganier, J. and L. Eggert, “Host Identity Protocol (HIP) Rendezvous Extension,” June 2006.)) and "HIP Extensions for the Traversal of Network Address Translators" ([I-D.ietf-hip-nat-traversal] (Komu, M., Henderson, T., Tschofenig, H., Melen, J., and A. Keranen, “Basic HIP Extensions for Traversal of Network Address Translators,” October 2009.)).

- Other HIP applications would also benefit. For example, the HIP-HOP proposal in the P2PSIP working group ([I-D.matthews-p2psip-hip-hop] (Cooper, E., “A Distributed Transport Function in P2PSIP using HIP for Multi-Hop Overlay Routing,” June 2007.))includes a mechanism that allows intermediate peers to route HIP packets between peers that does not have an active direct connection (for various reasons, including but not limited to NAT traversal issues).

Although a HIP-level relay is attractive for these mechanisms, the HIP packet checksum is recomputed at each relay, because the source and

destination IP addresses are currently included in the pseudo-header covered by the HIP checksum.
In addition to the checksum calculation overhead at each relay, intermediate recalculation means that checksum protection only covers the part of the path to the next HIP-aware device that recalculates the checksum.
We also note that excluding the source and destination IP addresses from the HIP checksum calculation would improve the ability of HIP connections to survive interface selection changes, would better accommodate multihoming, and would better detect checksum failures in scenarios where intermediate nodes forward at the HIP layer. For example, in [I-D.ietf-hip-nat-traversal] (Komu, M., Henderson, T., Tschofenig, H., Melen, J., and A. Keranen, “Basic HIP Extensions for Traversal of Network Address Translators,” October 2009.), checksums are recalculated when packets traverse HIP-aware NATs (so checksums do not protect the packets over an entire end-to-end path) and checksums are set to 0 when packets traverse HIP-unaware NATs (so that checksums are disabled completely).

---

4. Changes to Checksum Coverage

TOC

This specification recommends removing pseudo-header fields from the checksum calculation, so that the checksum does not need to be recalculated when pseudo-header field values change. Most important (and probably most controversial) is the removal of source and destination IP addresses from the checksum calculation.
The only bytes included in the checksum calculation are the HIP header and HIP payload bytes.
The resulting checksum does protect the HIP header and HIP payload, but does not protect the pseudo-header fields associated with the HIP header and HIP payload. The resulting checksum is "end-to-end" - it protects the HIP header and HIP payload for packets that traverse HIP relays, where the current checksum calculation does not - it is recalculated at each HIP relay, so errors at intermediate HIP relays may not be caught if the checksum was recalculated after the error was introduced at the HIP relay.
We believe that removal of the source and destination IP addresses from the transport checksum calculation is a reasonable change to the HIP protocol. If these fields are exluded from the transport checksum calculation, we see no reason to include other IP-level fields as part of a transport pseudo-header in the checksum calculation.

TOC

5. Proposed HIP Checksum Coverage

This proposal uses the same checksum algorithm that is used in SCTP, except that no pseudo-header fields are prepended to the HIP header and payload before the checksum is calculated.
Note that the current HIP base specification uses a 16-bit checksum field. This checksum field must be expanded to 32 bits, in order to accommodate the crc32c checksum (also used by SCTP). Also note that a reserved field is also added to preserve 8-byte alignment, as required in [RFC2460] (Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," December 1998.) to ensure alignment of subsequent IPv6 extension headers.
The resulting modified header format is as follows:

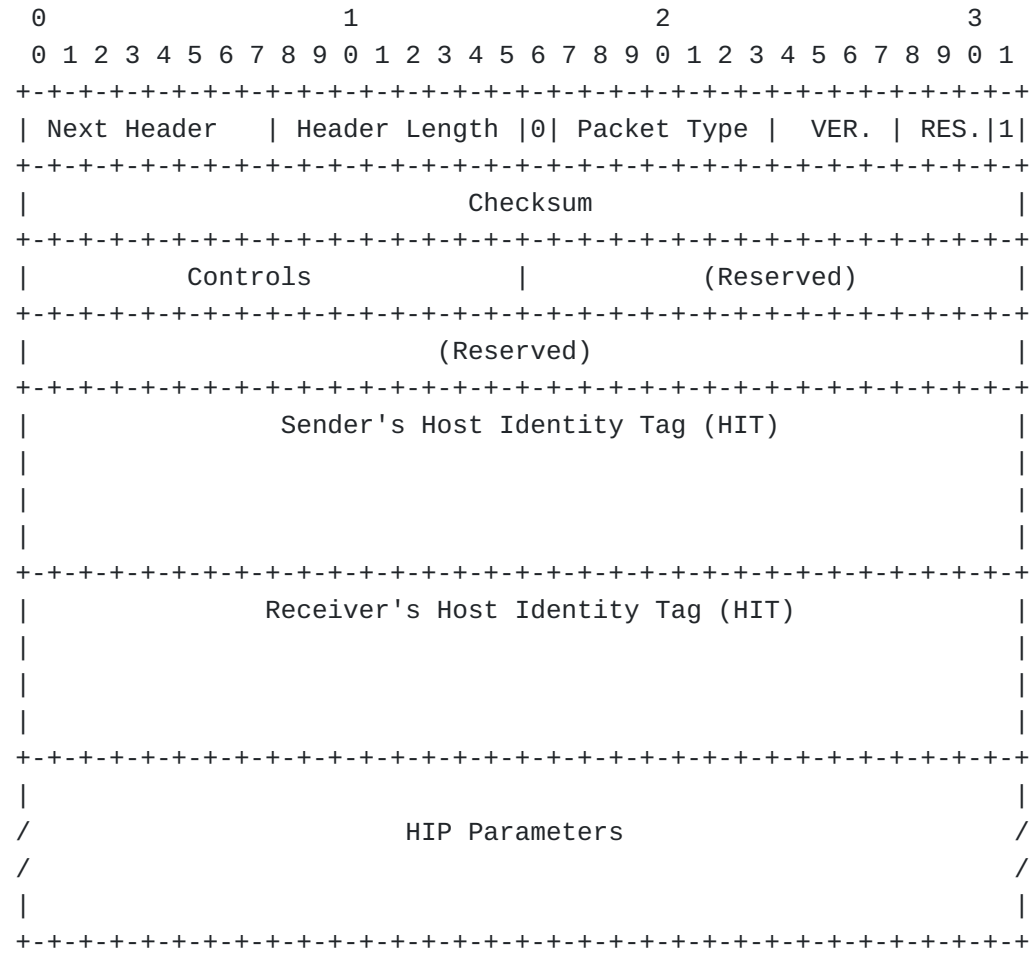

The following text is taken from [RFC4960] (Stewart, R., "Stream Control Transmission Protocol," September 2007.), with modifications (reflecting its use for HIP):
When sending a HIP packet, the endpoint MUST strengthen the data integrity of the transmission by including the crc32c checksum value calculated on the packet, as described below.
After the packet is constructed, the transmitter shall:

1. Initialize the checksum field to 0's.

2. Calculate the crc32c checksum of the HIP header and HIP payload. Refer to appendix B of [RFC4960] (Stewart, R., "Stream Control Transmission Protocol," September 2007.) for details of the crc32c algorithm. And,

3. Put the resultant value into the checksum field, and leave the rest of the bits unchanged.

When a HIP packet is received, the receiver MUST first check the crc32c checksum:

1. Store the received crc32c checksum value aside,

2. Replace the 32 bits of the checksum field in the received HIP packet with all '0's and calculate an crc32c checksum value of the HIP header and payload. And,

3. Verify that the calculated crc32c checksum is the same as the received crc32c checksum. If not, the receiver MUST treat the packet as an invalid HIP packet.

The default procedure for handling invalid HIP packets is to silently discard them.

---

6. Proposed HIP Version Change

TOC

Because this change is not backward-compatible with current HIP specifications ([I-D.ietf-hip-base] (Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson, "Host Identity Protocol," October 2007.)) or with conformant implementations to this specification, we recommend changing the HIP version number from 1 to 2 for packets processed using this specification.
Although it is certainly possible to make this change backward-compatible by adding a version negotiation mechanism and continuing to use HIP Version 1 with hosts that do not respond to the negotiation mechanism, we believe that making this change to HIP at this time, without providing such a negotiation mechanism, is appropriate because

1. The base HIP protocol specification is still an Internet Draft, as of this writing,

2. The base HIP protocol specification is currently targeted at the Experimental track, so that the usual RFC 2026

considerations about incompatible changes to Standards-track protocols should not apply, and

3. The change is small and localized to setting/verifying the HIP version number itself, plus computing and verifying the HIP checksum.

---

7. Security Considerations

TOC

The following security considerations are in addition to the security considerations identified in the base HIP protocol specification ([I-D.ietf-hip-base] (Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson, “Host Identity Protocol,” October 2007.)).
This change to the HIP checksum coverage does make source and destination IP addresses less "visible" to a receiving host, but this is judged to be an acceptable risk, for these reasons:

1. Off-path attackers must provide two 128-bit numbers (the source and destination HITs) that the host under attack will recognize as legitimate HITs. If the attacker sends packets with HITs that the host under attack does not recognize, these packets are simply dropped as invalid packets - a MUST in [I-D.ietf-hip-base] (Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson, “Host Identity Protocol,” October 2007.). Given that each of the two numbers is self-cerifying because it is the hash of a public key, this attack is seen as computationally hard.

2. On-path attackers who are able to learn valid source and destination HITs, in order to forge packets that might be used for denial of service attacks, etc. can already do much worse than forge packets - they can simply drop packets, or execute denial of service attacks on intermediate devices that would forward legitimate packets in normal operation.

---

8. IANA Considerations

TOC

The current HIP base protocol specification ([I-D.ietf-hip-base] (Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson, “Host Identity Protocol,” October 2007.)) includes as part of its IANA considerations a request that IANA create a new namespace for HIP Version Numbers, and add an entry for HIP Version 1. This specification

defines HIP Version 2, so we also request an entry in this namespace for HIP Version 2, as described in this specification.

9. References

TOC

9.1. Normative References

TOC

[I-D.ietf-hip-base]	Moskowitz, R., Nikander, P., Jokela, P., and T. Henderson, “Host Identity Protocol,” draft-ietf-hip-base-10 (work in progress), October 2007 (TXT).
[RFC2119]	Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” BCP 14, RFC 2119, March 1997 (TXT, HTML, XML).
[RFC4960]	Stewart, R., “Stream Control Transmission Protocol,” RFC 4960, September 2007 (TXT).

9.2. Informative References

TOC

[I-D.ietf-hip-nat-traversal]	Komu, M., Henderson, T., Tschofenig, H., Melen, J., and A. Keranen, “Basic HIP Extensions for Traversal of Network Address Translators,” draft-ietf-hip-nat-traversal-09 (work in progress), October 2009 (TXT).
[I-D.ietf-hip-rvs]	Laganier, J. and L. Eggert, “Host Identity Protocol (HIP) Rendezvous Extension,” draft-ietf-hip-rvs-05 (work in progress), June 2006 (TXT).
[I-D.matthews-p2psip-hip-hop]	Cooper, E., “A Distributed Transport Function in P2PSIP using HIP for Multi-Hop Overlay Routing,” draft-matthews-p2psip-hip-hop-00 (work in progress), June 2007 (TXT).
[RFC0768]	Postel, J., “User Datagram Protocol,” STD 6, RFC 768, August 1980 (TXT).
[RFC2460]	Deering, S. and R. Hinden, “Internet Protocol, Version 6 (IPv6) Specification,” RFC 2460, December 1998 (TXT, HTML, XML).

Authors' Addresses

TOC

	XingFeng Jiang
	Huawei Technologies

	Huihong Mansion,No.91 Baixia Rd
	Nanjing, Jiangsu 210001
	P. R. China
Phone:	+86(25)84565462
Email:	jiang.x.f@huawei.com
	Philip Matthews
	Avaya
	100 Innovation Drive
	Ottawa, Ontario K2K 3G7
	Canada
Phone:	+1 613 592 4343 x224
Email:	philip_matthews@magma.ca
	Spencer Dawkins (editor)
	Huawei Technologies (USA)
	1547 Rivercrest Blvd.
	Allen, TX 75002
	USA
Phone:	+1 214 755 3870
Email:	spencer@mcsr-labs.org

Full Copyright Statement

TOC

Copyright © The IETF Trust (2007).
This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.
This document and the information contained herein are provided on an “AS IS” basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights 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; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.