Network Working Group
Internet-Draft
Intended status: Experimental
Expires: July 14, 2014

B. Williams
Akamai, Inc.
M. Boucadair
France Telecom
D. Wing
Cisco Systems, Inc.
January 10, 2014

Experimental Option for TCP Host Identification
draft-williams-exp-tcp-host-id-opt-00

Abstract

Recent IETF proposals have identified benefits to more distinctly identifying the hosts that are hidden behind a shared address/prefix sharing device or application-layer proxy. Analysis indicates that the use of a TCP option for this purpose can be successfully applied to a broad range of use cases. This document describes a common experimental TCP option format for host identification.

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.

1. Introduction

A broad range of issues associated with address sharing have been well documented in [RFC6269] and [I-D.boucadair-intarea-host-identifier-scenarios]. In addition, [RFC6967] provides analysis of various solutions to the problem of revealing the sending hosts's identifier (HOST_ID) information to the receiver, which indicates that a solution using a TCP [RFC0793] option for this purpose can be successfully applied to a broad range of use cases with limited performance impact.

Multiple recent Internet Drafts define TCP options for the purpose of host identification: [I-D.wing-nat-reveal-option], [I-D.abdo-hostid-tcpopt-implementation], and [I-D.williams-overlaypath-ip-tcp-rfc]. This document defines a common TCP option format to meet the needs of all three of the above proposals. The option defined in this document uses the TCP experimental option codepoint sharing mechanism defined in [RFC6994] and is intended to allow validation of this common option format in order to conduct more experimental work that will complement the experiment results already documented in [I-D.abdo-hostid-tcpopt-implementation].

Section 5 (Section 5) of this document discusses compatibility between this new TCP option and existing commonly deployed TCP options.

2. Terminology

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].

3. Option Format

When used for host identification, the TCP experimental option has the following format and content.

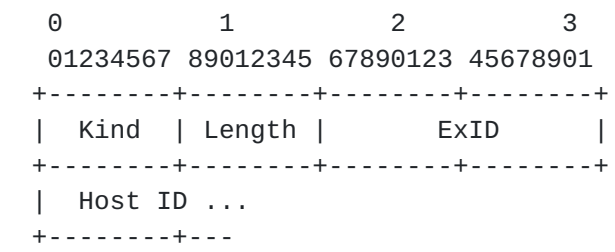

Kind: The option kind value is 253

Length: The length of the option is variable, based on the required size of the host identifier (e.g. a 2 octet host ID will require a length of 6, while a 4 octet host ID will require a length of 8).

ExID: The experiment ID value is 0x0348 (840).

Host ID: The host identifier is an application dependent value with an interpretation agreed upon by the sender and the receiver.

When multiple host identifiers are required (e.g. [I-D.williams-overlaypath-ip-tcp-rfc] defines an option that provides multiple IPv4 addresses, and [I-D.abdo-hostid-tcpopt-implementation] defines an option that may provide both an address and a port), the HOST_ID option is included multiple times within the packet, once for each identifier. While this approach significantly increases option space utilization when multiple identifiers are required, cases where only a single identifier is required are more common and thus it is beneficial to optimize for those cases.

4. Option Use

Intermediary devices (e.g. address sharing device) SHOULD be configurable to enable including the HOST_ID TCP option. These devices MUST be configured with the type of information to populate the HOST_ID TCP option (e.g. certain bits of the source IPv6 address, the full source IPv6 address, certain bits of the source IPv4 address, the full source IPv4 address, the source port number, etc.).

The device may be configured to include multiple identifiers (e.g. both a source IP address and a source port number). In such case, the device MUST insert two instances of the HOST_ID option, each of which contains the appropriate information. Note, there is no need to signal the semantic of the included data as this specification assumes the service is aware of that information by out of band means (e.g. both the service and the address sharing device are managed by the same administrative entity).

When an intermediary device is configured to include the HOST_ID option, it MUST include the HOST_ID TCP option in SYN messages. In addition, an intermediary device and a receiving end device MAY be configurable to allow inclusion of the HOST_ID TCP option in additional messages in order to support the use of SYN cookies. For example:

- The HOST_ID option from the initial SYN might be included in the SYN/ACK message when a SYN cookie is being sent in order to echo the HOST_ID value back to the intermediary device.

- The HOST_ID option might be included in ACK messages that contain no data.

- The HOST_ID option might be included in all ACK messages until return messages from the receiver positively indicate that an ACK has been received (e.g. the return messages either includes or acknowledges data).

The option SHOULD NOT be included in packets if the resulting packet would require local fragmentation. The option MUST NOT be include in packets when there is not enough space for at least one valid identifier of the configured type.

The device MUST be configured with the behavior to follow when a HOST_ID TCP option is already present in the message:

- If the device is configured to strip any existing HOST_ID TCP option, it MUST remove any occurrence of the HOST_ID in a received TCP message.

- If the device is configured to strip any existing HOST_ID TCP option and insert a local HOST_ID TCP Option, it MUST remove any occurrence of the HOST_ID in a received TCP message and then MUST include a local HOST_ID TCP option.

- The device may be configured to maintain any existing HOST_ID TCP option(s) in the received message, the device MUST NOT remove those instances of the option. Furthermore, it MUST add a new HOST_ID TCP option while preserving the order of appearance in the message. In particular, the local HOST_ID TCP option MUST appear as the last occurrence of the HOST_ID TCP option in the message.

5. Interaction with Other TCP Options

This section details how the HOST_ID option functions in conjunction with other TCP options.

5.1. Option Space

TCP provides for a maximum of 40 octets for TCP options. As discussed in Appendix A of Multipath TCP (MPTCP) [RFC6824], a typical SYN from modern, popular operating systems contain several TCP options (MSS, window scale, SACK permitted, and timestamp) which consume 19-24 octets depending on word alignment of the options. The initial SYN from a multipath TCP client would consume an additional 12 octets.

To save option space, the intermediate device adding the HOST_ID Option can break word-alignment of the TCP options, ensuring 40-19=21 octets (without MPTCP) or 40-19-12=9 octets (with MPTCP) are available for the HOST_ID option and its value. If, however, the intermediate device preserves word alignment (perhaps for compatibility with TCP servers that need word alignment), the intermediate device is left with less space: 40-24=16 octets (without MPTCP) or 40-24-12=4 octets (with MPTCP).

HOST_ID needs at least 6 octets to be useful, so 9-21 octets are sufficient for many scenarios that benefit from HOST_ID. However, 4 octets are not enough space for the HOST_ID option. Thus, a TCP SYN containing all the typical TCP options (MSS, window Scale, SACK permitted, timestamp), and also containing multipath capable or multipath join), and also being word aligned, has insufficient space to also accommodate HOST_ID. This means something has to give. The choices are to avoid word alignment in that case (freeing 5 octets), remove a TCP option from the original TCP SYN, or avoid adding the HOST_ID option. We expect to learn from deployment experience during the experiment which of these options, or a combination of these options, is best.

5.2. Authentication Option (TCP-AO)

The TCP-AO option [RFC5925] is incompatible with an intermediate device adding the HOST_ID option because TCP-AO provides integrity protection of the TCP SYN, including TCP options. However, TCP-AO is already incompatible with address sharing, because TCP-AO provides integrity protection of the source IP address. So the incompatibility with TCP-AO is not a problem in practice.

6. Security Considerations

Security (including privacy) considerations common to all HOST_ID solutions are discussed in [RFC6967]. These considerations should be taken into account.

7. Privacy Considerations

Sending a TCP SYN across the public internet necessarily discloses the public IP address of the sending host. When an intermediate address sharing device is deployed on the public internet (see [I-D.boucadair-intarea-host-identifier-scenarios] for examples), anonymity of the hosts using the device will be increased, with hosts represented by multiple source IP addresses on the ingress side of the device using a single source IP address on the egress side. The HOST_ID TCP option removes that increased anonymity, taking information that was already visible in TCP packets on the public internet on the ingress side of the address sharing device and making it available on the egress side of the device as well. In some cases, an explicit purpose of the address sharing device is anonymity, in which case use of the HOST_ID TCP option would be incompatible with the purpose of the device.

Use of the HOST_ID TCP option described here should follow the recommendations laid out in [RFC6967]. In particular:

- The HOST_ID option SHOULD NOT be used to provide client geographic or network location information that was not publicly visible in IP packets for the TCP flows processed by the inserting host. For example, the client's IP address MAY be used as the HOST_ID option value, but any geographic or network location information derived from the client's IP address SHOULD NOT be used as the HOST_ID value.

- The HOST_ID option MAY provide differentiating information that is locally unique such that individual TCP flows processed by the inserting host can be reliably identified. The HOST_ID option SHOULD NOT provide client identification information that was not publicly visible in IP packets for the TCP flows processed by the inserting host.

- The HOST_ID option SHOULD be stripped from IP packets traversing middle boxes that provide network-based anonymity services.

8. IANA Considerations

This document specifies a new TCP option that uses the shared experimental options format [RFC6994], with ExID=0x0348 (840) in network-standard byte order. This ExID has already been registered with IANA.

9. References

9.1. Normative References

[RFC0793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. Informative References

[I-D.abdo-hostid-tcpopt-implementation] Abdo, E., Boucadair, M., and J. Queiroz, "HOST_ID TCP Options: Implementation & Preliminary Test Results", draft-abdo-hostid-tcpopt-implementation-03 (work in progress), July 2012.

[I-D.boucadair-intarea-host-identifier-scenarios] Boucadair, M., Binet, D., Durel, S., Chatras, B., Reddy, T., and B. Williams, "Host Identification: Use Cases", draft-boucadair-intarea-host-identifier-scenarios-03 (work in progress), March 2013.

[I-D.williams-overlaypath-ip-tcp-rfc] Williams, B., "Overlay Path Option for IP and TCP", draft-williams-overlaypath-ip-tcp-rfc-04 (work in progress), June 2013.

[I-D.wing-nat-reveal-option] Yourtchenko, A. and D. Wing, "Revealing hosts sharing an IP address using TCP option", draft-wing-nat-reveal-option-03 (work in progress), December 2011.

[RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, June 2010.

[RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, June 2011.

[RFC6824] Ford, A., Raiciu, C., Handley, M., and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses", RFC 6824, January 2013.

[RFC6967] Boucadair, M., Touch, J., Levis, P., and R. Penno, "Analysis of Potential Solutions for Revealing a Host Identifier (HOST_ID) in Shared Address Deployments", RFC 6967, June 2013.

[RFC6994] Touch, J., "Shared Use of Experimental TCP Options", RFC 6994, August 2013.

Authors' Addresses

Brandon Williams
Akamai, Inc.
8 Cambridge Center
Cambridge, MA 02142
USA

Email: brandon.williams@akamai.com

Mohamed Boucadair
France Telecom
Rennes, 35000
Fance

Email: mohamed.boucadair@orange.com

Dan Wing
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
USA

Email: dwing@cisco.com