Network Working Group	T. Hain
Internet-Draft	Cisco Systems
Intended status: Standards Track	July 10, 2010
Expires: January 11, 2011	

TOC

IPv6 Firewall Routing Header
draft-hain-ipv6-fwrh-03

Abstract

This document specifies a routing header for use by firewalls to enforce routing symmetry.
The draft is being discussed on the ipv6@ietf.org list.

Legal

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction
2. Terminology
3. Mechanisms
    3.1. Overview
    3.2. Firewall Routing Header Option
    3.3. Routing Header Required ICMP Message
    3.4. ***** This needs to be a separate Doc ***** ReturnPathForwarding ICMP Error Message
4. Examples
5. IANA Considerations
6. Security Considerations
7. Acknowledgements
8. Pending comments
9. References
    9.1. Normative References
    9.2. Informative References
§ Author's Address

---

1. Introduction

TOC

With the deprecation of RH0 [RFC5095] (Abley, J., “Deprecation of Type 0 Routing Headers in IPv6,” December 2007.)the ability of a node to influence traffic to traverse the same firewall in opposing directions was eliminated. Operation of stateful firewalls requires path symmetry at least through that device. The likelihood of asymmetry is particularly true on the Internet side, but is also a problem when the exit path changes on the private side during an established packet exchange. This document targets the specific use case of symmetric firewall selection, and then defines an IPv6 Firewall Routing Header and associated IPv6 ICMP error messages. Other use cases may take advantage of these mechanisms, but are explicitly out of scope for the development and definition here. Also, this option is not trying to solve every possible topology of firewall deployments, or source of asymmetry, but is should be useful for the most common existing deployments.

While this Firewall Routing Header mechanism allows for changes in the firewalls in use during a packet exchange, any mechanisms that would be used to pass state between disperse firewalls is out of scope here. If those state passing mechanisms exist, it will be possible for an end-to-end packet exchange to persist during a routing shift, even over vast geographic path changes. Since the ability to force routing through a single intermediary can be used by attackers to receive traffic they otherwise would not, the ICMP error messages used here should be integrity protected, and in any event should be dropped if originating outside of a policy domain. The method that a node would use to verify that signature is out of scope here.

---

2. Terminology

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 RFC 2119 (Bradner, S., “Key words for use in RFCs to Indicate Requirement Levels,” March 1997.) [RFC2119].

---

3. Mechanisms

TOC

3.1. Overview

TOC

To deal with the asymmetry issue, a function specific IPv6 routing header is defined, the Firewall Routing Header (FWRH). This routing header is used to pass information to a correspondent, so that return packets are routed through a firewall that is maintaining state for this transaction. In the case where an organization deploys multiple firewalls from different vendors in series to reduce vendor specific issues, this option would be related to the most public facing firewall of the set. A flag (Origin Flag) will be used to indicate which instance this is in the case where there are firewalls on each end (note: this assumes that a firewall would allow an initial inbound packet, then enforce subsequent packets through itself). There can be at most 2 (one with each value of the Origin Flag) FWRH's in any given extension header chain.

An originating node is probably unaware of the presence of a firewall on its path to any given destination, so it will likely send an initial packet without the FWRH option. (It may have cached the value received in a prior ICMP error message for non-local prefixes, and if so can optimistically minimize the chance for an initial error by including the cached value in the optional FWRH.) On receipt of any ICMP Routing Header Required message, the origin node will extract the return-via address from the ICMP, optionally verify any signature that may be present, then cache it for use over the lifetime of the current packet exchange, and optionally for use with other non-local prefixes. The packet which generated the error is reconstructed with the appropriate FWRH option, and resent. If the FWRH option with the appropriate Origin Flag value is missing when a packet is being sent outbound (from the perspective of the firewall - private toward public), or if the return-via address in the FWRH option does not match any address on the firewall that is processing the packet, an ICMP error (RHR type ???) is returned to the source address of the packet, indicating the value that is expected to result in return-path symmetry through an interface on this specific firewall. That ICMP error SHOULD be signed, and if so the ICMP packet indicates which signing method was used. Outbound packets with the matching Origin Flag and return-via address will receive normal firewall handling before forwarding. If the source address prefix does not match any prefix that being exchanged in routing protocols with the next hop, an RPF-error (type ???+1) ICMP message should be returned.

The correspondent node will extract the return-via address from a received FWRH option with an Origin Flag that is opposite the one it will use when sending, and cache it for use related to this specific packet exchange. When constructing packets related to this specific exchange; after all normal processing is complete, the ultimate destination address will be swapped into the FWRH with the appropriate Origin Flag, and the previously cached value from any received FWRH with the opposing Origin Flag will be swapped as the initial destination address of the packet.

Inbound packets to the public side of a firewall will be addressed with it as the destination address of the base IPv6 header, and the FWRH with the Origin flag matching existing state will contain the address of the ultimate destination. The contents of the FWRH with the matching Origin flag will be swapped with the destination addresses in the packet, before normal firewall processing and forwarding.

3.2. Firewall Routing Header Option

The Firewall Routing header is used by an IPv6 source to list a single intermediate node to be "visited" by returning packets. The Firewall Routing header is identified by a Next Header value of 43 in the immediately preceding header, and has the following format:

```
  0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |  Next Header  |  Hdr Ext Len  |  Routing Type |     Flags     |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                             Reserved                          |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 +                             Address                           +
 |                         return-via or                         |
 +                       ultimate destination                    +
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Next Header 8-bit selector.
Identifies the type of header immediately following the Routing header. Uses the same values as the IPv4 Protocol field   [RFC-1700 et seq.].

Hdr Ext Len 8-bit unsigned integer. Length of the Routing header in 8-octet units, not including the first 8 octets.

Routing Type 8-bit identifier of this particular Routing header variant.

Origin Flag The Origin Flag (Flags - bit 0) is used to indicate the direction that corresponds to this instance of the FWRH. A value of 1 indicates that a firewall exists in the initial packet direction (for TCP this is the SYN), while a value of 0 indicates that a firewall exists in the response direction (for TCP this is the SYN-ACK).

Reserved   32-bit reserved field.
Initialized to zero for transmission; ignored on reception.

Address   Direction/location dependent address, where from the origin node toward a correspondent it contains the address of the firewall to be returned through and from the correspondent to the intermediate firewall contains the ultimate destination address that the firewall will deliver the packet to.

**** Within an enterprise network both FWRH options might contain the same address if there was an audit function, or traffic engineering reason to route both directions through the same mid-point.

A FWRH option SHOULD NOT contain an address that matches either the source or destination addresses of the packet. If there are two FWRH options present: Their Origin Flag values MUST be different Both options SHOULD NOT contain the same address

---

3.3. Routing Header Required ICMP Message

TOC

Informing the originating endpoint that it needs to insert a routing header is accomplished with a specific ICMP error - 'Routing Header Required' (type ???). This ICMP error may optionally be signed to mitigate a man-in-the-middle attack vector which could be used to route all return-path traffic through an attacker's node.

```
  0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |     Type      |     Code      |          Checksum             |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                            Reserved                           |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 +                                                               +
 |                     Return-Via  IPv6 address                  |
 +                                                               +
 |                                                               |
 +                                                               +
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |     MAC-type      |             MAC-length                    |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 +                                                               +
 |                 Message Authentication Code                   |
 +                                                               +
 |                                                               |
 +                                                               +
```

IPv6 Fields:

Destination Address

Copied from the Source Address field of the invoking packet.

ICMPv6 Fields:

Type ???? --- 5

Code 0 - unsigned
1 - signed

Reserved 32-bit reserved field. Initialized to zero for transmission; ignored on reception.

MAC-type 0 - unused
1 - HMAC-SHA1
2 - AES-CMAC
...

MAC-length Length of the message authentication option

MAC   Value of the message authenticator

The optional authentication covers the source and destination address of this specific packet exchange, plus the return-via address.

3.4. ***** This needs to be a separate Doc ***** ReturnPathForwarding ICMP Error Message

TOC

Informing the origin that the source prefix is not appropriate for the next hop as a symmetric return path.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Type      |     Code      |          Checksum             |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Prefix Length                         |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                                                               +
|                 Source Prefix for this Path                   |
+                                                               +
|                                                               |
+                                                               +
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     MAC-type      |             MAC-length                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
+                                                               +
|                 Message Authentication Code                   |
+                                                               +
|                                                               |
+                                                               +
```

IPv6 Fields:

Destination Address

Copied from the Source Address field of the invoking packet.

ICMPv6 Fields:

Type ???? --- 6

Code 0 - unsigned
1 - signed

Reserved 32-bit reserved field. Initialized to zero for transmission; ignored on reception.

MAC-type 0 - unused
1 - HMAC-SHA1
2 - AES-CMAC
...

MAC-length Length of the message authentication option

MAC   Value of the message authenticator

The optional authentication covers the source and destination address of this specific packet exchange, plus the return-via address.

4. Examples

TOC

Example 1 - Single firewall connecting a private network to the Internet.

```
            A           FW1              B
            |__Private__|  |__Internet__|

     Initial packet exchange sequence -
     A   ->>         DST = B
     FW1 ICMP(5) <<- DST = A : SRC = FW1(P) : value - FW1(I)
     A   ->>         DST = B : FWRH(OF=1) = FW1(I)
     B   <<-         DST = FW1(I) : FWRH(OF=1) = A
     FW1 <<-         DST = A : FWRH(OF=1) = FW1(I)
     A   ->>         DST = B : FWRH(OF=1) = FW1(I)
     B   <<-         DST = FW1(I) : FWRH(OF=1) = A
     FW1 <<-         DST = A : FWRH(OF=1) = FW1(I)
                             ...
```

Figure 1

In this case as node A attempts to connect to node B, FW1 rejects the first attempt with an error message informing that the FWRH is required. The subsequent packet including the FWRH with Origin Flag set is forwarded by FW1 toward B. B sends the return packet to FW1, where the DST is swapped with the locator for A in the FWRH with the Origin Flag set.

```
             A            FW1             B
             |__Private__|  |__Internet__|
                         |                |
                         FW2 _ Internet _|

      Initial packet exchange sequence -
      A   ->>         DST = B
      FW1 ICMP(5) <<- DST = A : SRC = FW1(P) : value - FW1(I)
      A   ->>         DST = B : FWRH(OF=1) = FW1(I)
      B   <<-         DST = FW1(I) : FWRH(OF=1) = A
      FW1 <<-         DST = A : FWRH(OF=1) = FW1(I)
              --- FW1 dies
      A   ->>         DST = B : FWRH(OF=1) = FW1(I)
      FW2 ICMP(5) <<- DST = A : SRC = FW2(P) : value - FW2(I)
      FW2 ICMP(6) <<- DST = A : SRC = FW2(P) : value - /48 FW2(I)
      A   ->>         DST = B : FWRH(OF=1) = FW2(I)
      B   <<-         DST = FW2(I) : FWRH(OF=1) = A
      FW2 <<-         DST = A : FWRH(OF=1) = FW1(I)
                               ...
```

Figure 2

This case is the same as above, except there is an alternate firewall available to A. At some point after the connection is established, FW1 dies, and routing redirects packets to B through FW2. FW2 has acquired state from FW1, so the connection between A & B does not have to be reset, but FW2 still rejects the next packet with an error message informing that the FWRH does not have the public address of FW2. The subsequent packet including the FWRH with Origin Flag set is forwarded by FW2 toward B. B sends the return packet to FW2, where the DST is swapped with the locator for A in the FWRH with the Origin Flag set.

```
            A           FW1              FW2           B
            |__Private__|  |__Internet__| |__Private__|

     Initial packet exchange sequence -
     A    ->>          DST = B
     FW1 ICMP(5) <<- DST = A : SRC = FW1(P) : value - FW1(I)
     A    ->>          DST = B : FWRH(OF=1) = FW1(I)
     FW2 ->>           (presumes that FW2 allows initial pkt without state)
     B    <<-          DST = FW1(I) : FWRH(OF=1) = A
     FW2 ICMP(5) ->> DST = B : SRC = FW2(P) : value - FW2(I)
     B    <<-          DST = FW1(I) : FWRH(OF=1) = A : FWRH(OF=0) = FW2(I)
     FW1 <<-           DST = A : FWRH(OF=1) = FW1(I) : FWRH(OF=0) = FW2(I)
     A    ->>          DST= FW2(I) : FWRH(OF=1) = FW1(I) : FWRH(OF=0) = B
     FW2 ->>           DST = B : FWRH(OF=1) = FW1(I) : FWRH(OF=0) = FW2(I)
                                ...
```

Figure 3

In this case there are firewalls at each end, and both require a FWRH. The value of the Origin Flag identifies which FWRH option is associated with each firewall. Note that before forwarding to the private side of each firewall, the DST & FWRH(OF=1) was swapped at FW1, while DST & FWRH(OF=0) was swapped at FW2.

5. IANA Considerations

TOC

This specification registers a Routing Header type & two ICMP message types

6. Security Considerations

TOC

A routing header is used to cause packets to traverse a specific node, and if used maliciously would allow an attacker to see all packets in an exchange. The risk of this attack is minimized by filtering out any ICMP Routing Header Required and ICMP RPF-error messages that originate outside the policy domain, and/or signing & verifying those ICMP error messages when generated internally.

TOC

7. Acknowledgements

The need to resolve routing symmetry for firewalls was initially championed by William Dixon, and discussed with many attendees at IETF 74.

8. Pending comments

TOC

It has been suggested the Origin Flag model will fail in simultaneous-open situations. Recommendation to change the OF to indicate src < dst. If that was done as a rule, there wouldn't need to be a flag.

9. References

TOC

9.1. Normative References

TOC

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

9.2. Informative References

TOC

[RFC5095]	Abley, J., "Deprecation of Type 0 Routing Headers in IPv6," RFC 5095, December 2007 (TXT, HTML, XML).

Author's Address

TOC

	Tony Hain
	Cisco Systems
	500 108th Ave NE
	Bellevue, WA 98004
	USA
Phone:	+1 425 468-1061
Email:	alh-ietf@tndh.net