Internet Engineering Task Force (IETF)
Request for Comments: 7676
Category: Standards Track
ISSN: 2070-1721

C. Pignataro
Cisco Systems
R. Bonica
Juniper Networks
S. Krishnan
Ericsson

October 2015

IPv6 Support for Generic Routing Encapsulation (GRE)

Abstract

Generic Routing Encapsulation (GRE) can be used to carry any network-layer payload protocol over any network-layer delivery protocol. Currently, GRE procedures are specified for IPv4, used as either the payload or delivery protocol. However, GRE procedures are not specified for IPv6.

This document specifies GRE procedures for IPv6, used as either the payload or delivery protocol.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7676.

Table of Contents

1. Introduction 3
1.1. Requirements Language 3
1.2. Terminology 3
2. GRE Header Fields 4
2.1. Checksum Present 4
3. IPv6 as GRE Payload 5
3.1. GRE Protocol Type Considerations 5
3.2. MTU Considerations 5
3.3. Fragmentation Considerations 5
4. IPv6 as GRE Delivery Protocol 6
4.1. Next Header Considerations 6
4.2. Checksum Considerations 6
4.3. MTU Considerations 8
5. Security Considerations 8
6. References 8
6.1. Normative References 8
6.2. Informative References 9
Acknowledgements 10
Authors' Addresses 11

1. Introduction

Generic Routing Encapsulation (GRE) [RFC2784] [RFC2890] can be used to carry any network-layer payload protocol over any network-layer delivery protocol. Currently, GRE procedures are specified for IPv4 [RFC791], used as either the payload or delivery protocol. However, GRE procedures are not specified for IPv6 [RFC2460].

This document specifies GRE procedures for IPv6, used as either the payload or delivery protocol. Like RFC 2784, this document describes how GRE has been implemented by several vendors.

1.1. Requirements Language

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].

1.2. Terminology

The following terms are used in this document:

- GRE delivery header: An IPv4 or IPv6 header whose source address represents the GRE ingress node and whose destination address represents the GRE egress node. The GRE delivery header encapsulates a GRE header.

- GRE header: The GRE protocol header. The GRE header is encapsulated in the GRE delivery header and encapsulates the GRE payload.

- GRE payload: A network-layer packet that is encapsulated by the GRE header.

- GRE overhead: The combined size of the GRE delivery header and the GRE header, measured in octets.

- Path MTU (PMTU): The minimum MTU of all the links in a path between a source node and a destination node. If the source and destination node are connected through Equal-Cost Multipath (ECMP), the PMTU is equal to the minimum link MTU of all links contributing to the multipath.

- Path MTU Discovery (PMTUD): A procedure for dynamically discovering the PMTU between two nodes on the Internet. PMTUD procedures for IPv6 are defined in [RFC1981].

- GRE MTU (GMTU): The maximum transmission unit, i.e., maximum packet size in octets, that can be conveyed over a GRE tunnel without fragmentation of any kind. The GMTU is equal to the PMTU associated with the path between the GRE ingress and the GRE egress, minus the GRE overhead.

2. GRE Header Fields

This document does not change the GRE header format or any behaviors specified by RFC 2784 or RFC 2890.

2.1. Checksum Present

The GRE ingress node SHOULD set the Checksum Present field in the GRE header to zero. However, implementations MAY support a configuration option that causes the GRE ingress node to set the Checksum Present field to one.

As per Section 2.2 of RFC 2784, the GRE egress node uses the Checksum Present field to calculate the length of the GRE header. If the Checksum Present field is set to one, the GRE egress node MUST use the GRE Checksum to verify the integrity of the GRE header and payload.

Setting the Checksum Present field to zero reduces the computational cost of GRE encapsulation and decapsulation. In many cases, the GRE Checksum is partially redundant with other checksums. For example:

- If the payload protocol is IPv4, the IPv4 header is protected by both the GRE Checksum and the IPv4 Checksum.

- If the payload carries TCP [RFC793], the TCP pseudo header, TCP header, and TCP payload are protected by both the GRE Checksum and TCP Checksum.

- If the payload carries UDP [RFC768], the UDP pseudo header, UDP header, and UDP payload are protected by both the GRE Checksum and UDP Checksum.

However, if the GRE Checksum Present field is set to zero, the GRE header is not protected by any checksum. Furthermore, depending on which of the above-mentioned conditions are true, selected portions of the GRE payload will not be protected by any checksum.

Network operators should evaluate risk factors in their networks and configure GRE ingress nodes appropriately.

3. IPv6 as GRE Payload

The following considerations apply to GRE tunnels that carry an IPv6 payload.

3.1. GRE Protocol Type Considerations

The Protocol Type field in the GRE header MUST be set to Ether Type [RFC7042] 0x86DD (IPv6).

3.2. MTU Considerations

A GRE tunnel MUST be able to carry a 1280-octet IPv6 packet from ingress to egress, without fragmenting the payload packet. All GRE tunnels with a GMTU of 1280 octets or greater satisfy this requirement. GRE tunnels that can fragment and reassemble delivery packets also satisfy this requirement, regardless of their GMTU. However, the ability to fragment and reassemble delivery packets is not a requirement of this specification. This specification requires only that GRE ingress nodes refrain from activating GRE tunnels that do not satisfy the above-mentioned requirement.

Before activating a GRE tunnel and periodically thereafter, the GRE ingress node MUST verify the tunnel's ability to carry a 1280-octet IPv6 payload packet from ingress to egress, without fragmenting the payload. Having executed those procedures, the GRE ingress node MUST activate or deactivate the tunnel accordingly.

Implementation details regarding the above-mentioned verification procedures are beyond the scope of this document. However, a GRE ingress node can verify tunnel capabilities by sending a 1280-octet IPv6 packet addressed to itself through the tunnel under test.

Many existing implementations [RFC7588] do not support the above-mentioned verification procedures. Unless deployed in environments where the GMTU is guaranteed to be greater than 1280, these implementations MUST be configured so that the GRE endpoints can fragment and reassemble the GRE delivery packet.

3.3. Fragmentation Considerations

When the GRE ingress receives an IPv6 payload packet whose length is less than or equal to the GMTU, it can encapsulate and forward the packet without fragmentation of any kind. In this case, the GRE ingress router MUST NOT fragment the payload or delivery packets.

When the GRE ingress receives an IPv6 payload packet whose length is greater than the GMTU, and the GMTU is greater than or equal to 1280 octets, the GRE ingress router MUST:

- discard the IPv6 payload packet
- send an ICMPv6 Packet Too Big (PTB) [RFC4443] message to the IPv6 payload packet source. The MTU field in the ICMPv6 PTB message is set to the GMTU.

When the GRE ingress receives an IPv6 payload packet whose length is greater than the GMTU, and the GMTU is less than 1280 octets, the GRE ingress router MUST:

- encapsulate the entire IPv6 packet in a single GRE header and IP delivery header
- fragment the delivery header, so that it can be reassembled by the GRE egress

4. IPv6 as GRE Delivery Protocol

The following considerations apply when the delivery protocol is IPv6.

4.1. Next Header Considerations

When the GRE delivery protocol is IPv6, the GRE header MAY immediately follow the GRE delivery header. Alternatively, IPv6 extension headers MAY be inserted between the GRE delivery header and the GRE header.

If the GRE header immediately follows the GRE delivery header, the Next Header field in the IPv6 header of the GRE delivery packet MUST be set to 47. If extension headers are inserted between the GRE delivery header and the GRE header, the Next Header field in the last IPv6 extension header MUST be set to 47.

4.2. Checksum Considerations

As stated in [RFC2784], the GRE header can contain a checksum. If present, the GRE header checksum can be used to detect corruption of the GRE header and GRE payload.

The GRE header checksum cannot be used to detect corruption of the IPv6 delivery header. Furthermore, the IPv6 delivery header does not contain a checksum of its own. Therefore, no available checksum can be used to detect corruption of the IPv6 delivery header.

In one failure scenario, the destination address in the IPv6 delivery header is corrupted. As a result, the IPv6 delivery packet is delivered to a node other than the intended GRE egress node. Depending upon the state and configuration of that node, it will either:

a. Drop the packet

b. Decapsulate the payload and forward it to its intended destination

c. Decapsulate the payload and forward it to a node other than its intended destination.

Behaviors a) and b) are acceptable. Behavior c) is not acceptable.

Behavior c) is possible only when the following conditions are true:

1. The intended GRE egress node is a Virtual Private Network (VPN) Provider Edge (PE) router.

2. The node to which the GRE delivery packet is mistakenly delivered is also a VPN PE router.

3. VPNs are attached to both of the above-mentioned nodes. At least two of these VPN's number hosts are from a non-unique (e.g., [RFC1918]) address space.

4. The intended GRE egress node maintains state that causes it to decapsulate the packet and forward the payload to its intended destination

5. The node to which the GRE delivery packet is mistakenly delivered maintains state that causes it to decapsulate the packet and forward the payload to an identically numbered host in another VPN.

While the failure scenario described above is extremely unlikely, a single misdelivered packet can adversely impact applications running on the node to which the packet is misdelivered. Furthermore, leaking packets across VPN boundaries also constitutes a security breach. The risk associated with behavior c) could be mitigated with end-to-end authentication of the payload.

Before deploying GRE over IPv6, network operators should consider the likelihood of behavior c) in their network. GRE over IPv6 MUST NOT be deployed other than where the network operator deems the risk associated with behavior c) to be acceptable.

4.3. MTU Considerations

By default, the GRE ingress node cannot fragment the IPv6 delivery header. However, implementations MAY support an optional configuration in which the GRE ingress node can fragment the IPv6 delivery header.

Also by default, the GRE egress node cannot reassemble the IPv6 delivery header. However, implementations MAY support an optional configuration in which the GRE egress node can reassemble the IPv6 delivery header.

5. Security Considerations

The Security Considerations section of [RFC4023] identifies threats encountered when MPLS is delivered over GRE. These threats apply to any GRE payload. As stated in RFC 4023, these various threats can be mitigated through options such as authenticating and/or encrypting the delivery packet using IPsec [RFC4301]. Alternatively, when the payload is IPv6, these threats can also be mitigated by authenticating and/or encrypting the payload using IPsec, instead of the delivery packet. Otherwise, the current specification introduces no security considerations beyond those mentioned in RFC 2784.

More generally, security considerations for IPv6 are discussed in [RFC4942]. Operational security for IPv6 is discussed in [OPSEC-V6], and security concerns for tunnels in general are discussed in [RFC6169].

6. References

6.1. Normative References

[RFC768] Postel, J., "User Datagram Protocol", STD 6, RFC 768, DOI 10.17487/RFC0768, August 1980, <http://www.rfc-editor.org/info/rfc768>.

[RFC791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <http://www.rfc-editor.org/info/rfc791>.

[RFC793] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, DOI 10.17487/RFC0793, September 1981, <http://www.rfc-editor.org/info/rfc793>.

[RFC1981] McCann, J., Deering, S., and J. Mogul, "Path MTU Discovery for IP version 6", RFC 1981, DOI 10.17487/RFC1981, August 1996, <http://www.rfc-editor.org/info/rfc1981>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, DOI 10.17487/RFC2460, December 1998, <http://www.rfc-editor.org/info/rfc2460>.

[RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", RFC 2784, DOI 10.17487/RFC2784, March 2000, <http://www.rfc-editor.org/info/rfc2784>.

[RFC2890] Dommety, G., "Key and Sequence Number Extensions to GRE", RFC 2890, DOI 10.17487/RFC2890, September 2000, <http://www.rfc-editor.org/info/rfc2890>.

[RFC4023] Worster, T., Rekhter, Y., and E. Rosen, Ed., "Encapsulating MPLS in IP or Generic Routing Encapsulation (GRE)", RFC 4023, DOI 10.17487/RFC4023, March 2005, <http://www.rfc-editor.org/info/rfc4023>.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, DOI 10.17487/RFC4301, December 2005, <http://www.rfc-editor.org/info/rfc4301>.

[RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, DOI 10.17487/RFC4443, March 2006, <http://www.rfc-editor.org/info/rfc4443>.

6.2. Informative References

[OPSEC-V6] Chittimaneni, K., Kaeo, M., and E. Vyncke, "Operational Security Considerations for IPv6 Networks", Work in Progress, draft-ietf-opsec-v6-07, September 2015.

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <http://www.rfc-editor.org/info/rfc1918>.

[RFC4942] Davies, E., Krishnan, S., and P. Savola, "IPv6 Transition/ Co-existence Security Considerations", RFC 4942, DOI 10.17487/RFC4942, September 2007, <http://www.rfc-editor.org/info/rfc4942>.

[RFC6169] Krishnan, S., Thaler, D., and J. Hoagland, "Security Concerns with IP Tunneling", RFC 6169, DOI 10.17487/RFC6169, April 2011, <http://www.rfc-editor.org/info/rfc6169>.

[RFC7042] Eastlake 3rd, D. and J. Abley, "IANA Considerations and IETF Protocol and Documentation Usage for IEEE 802 Parameters", BCP 141, RFC 7042, DOI 10.17487/RFC7042, October 2013, <http://www.rfc-editor.org/info/rfc7042>.

[RFC7588] Bonica, R., Pignataro, C., and J. Touch, "A Widely Deployed Solution to the Generic Routing Encapsulation (GRE) Fragmentation Problem", RFC 7588, DOI 10.17487/RFC7588, July 2015, <http://www.rfc-editor.org/info/rfc7588>.

Acknowledgements

The authors would like to thank Fred Baker, Stewart Bryant, Benoit Claise, Ben Campbell, Carlos Jesus Bernardos Cano, Spencer Dawkins, Dino Farinacci, David Farmer, Brian Haberman, Tom Herbert, Kathleen Moriarty, Fred Templin, Joe Touch, Andrew Yourtchenko, and Lucy Yong for their thorough review and useful comments.

Authors' Addresses

Carlos Pignataro
Cisco Systems
7200-12 Kit Creek Road
Research Triangle Park, North Carolina 27709
United States

Email: cpignata@cisco.com

Ron Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, Virginia
United States

Email: rbonica@juniper.net

Suresh Krishnan
Ericsson
8400 Decarie Blvd.
Town of Mount Royal, QC
Canada

Phone: +1 514 345 7900 x42871
Email: suresh.krishnan@ericsson.com