

Intarea Working Group
Internet-Draft
Updates: [2784](#) (if approved)
Intended status: Standards Track
Expires: September 28, 2015

C. Pignataro
Cisco Systems
R. Bonica
Juniper Networks
S. Krishnan
Ericsson
March 27, 2015

IPv6 Support for Generic Routing Encapsulation (GRE)
draft-ietf-intarea-gre-ipv6-04

Abstract

Generic Routing Encapsulation (GRE) can be used to carry any network-layer payload protocol over any network-layer delivery protocol. 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. It updates the GRE specification, [RFC 2784](#).

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](#)].

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 28, 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.	Terminology	3
2.	GRE Header Fields	3
2.1.	Checksum Present	3
2.2.	Protocol Type	3
3.	IPv6 as a GRE Payload	4
3.1.	MTU Considerations	4
4.	IPv6 as a GRE Delivery Protocol	5
4.1.	Checksum Considerations	5
4.2.	MTU Considerations	5
5.	IANA Considerations	6
6.	Security Considerations	6
7.	Acknowledgements	6
8.	Normative References	6
	Authors' Addresses	7

[1.](#) Introduction

Generic Routing Encapsulation (GRE) [[RFC2784](#)] [[RFC2890](#)] can be used to carry any network-layer payload protocol over any network-layer delivery protocol. GRE procedures are specified for IPv4 [[RFC0791](#)], 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. It updates [RFC 2784](#) [[RFC2784](#)]. Like [RFC 2784](#), this specification describes GRE how has been implemented by several vendors.

1.1. Terminology

The following terms are specific to GRE and are taken from [[RFC2784](#)]:

- o 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.
- o GRE header - the GRE protocol header. The GRE header is encapsulated in the GRE delivery header and encapsulates GRE payload.
- o GRE payload - a network layer packet that is encapsulated by the GRE header.

The following terms are specific MTU discovery:

- o 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.
- o 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](#)].

2. GRE Header Fields

This document does not change the GRE header format or any behaviors specified by [[RFC2784](#)] or [[RFC2890](#)].

2.1. Checksum Present

When the delivery protocol is IPv6, the GRE ingress router SHOULD set the Checksum Present field to zero. GRE egress routers MUST accept either a value of zero or one in this field. If the GRE egress router receives a value of one, it MUST use that information to calculate the GRE header length. It MUST also use the checksum to verify packet integrity.

2.2. Protocol Type

The Protocol Type field contains the protocol type of the payload packet. Protocol Types are defined in [[ETYPES](#)]. An implementation receiving a packet containing a Protocol Type which is not listed in [[ETYPES](#)] SHOULD discard the packet.

3. IPv6 as a GRE Payload

When the GRE payload is IPv6, the Protocol Type field in the GRE header MUST be set to 0x86DD.

3.1. MTU Considerations

The GRE ingress router maintains an estimate of the GRE MTU (GMTU). The GMTU is equal to the PMTU associated with the path between the GRE ingress and the GRE egress, minus the GRE overhead. The GRE overhead is the combined length of the GRE and IP delivery headers.

The GRE ingress router obtains a PMTU estimate using any of the following:

- o System defaults
- o Configuration
- o PMTUD

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:

- o discard the IPv6 payload packet
- o 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.

The GRE ingress router MUST support a configuration option that determines how the GRE ingress behaves when it receives an IPv6 payload packet whose length is greater than the GMTU, and the GMTU is less than 1280 octets. In its default configuration, the GRE ingress router MUST:

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

However, in an alternative configuration, the GRE ingress MAY:

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

4. IPv6 as a GRE Delivery Protocol

When the GRE delivery protocol is IPv6, the GRE header can 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.

4.1. Checksum Considerations

As stated in [[RFC2784](#)], the Checksum field contains the IP (one's complement) checksum sum of the all the 16 bit words in the GRE header and the payload packet. Therefore, the checksum does not ensure the integrity of the IPv6 delivery header.

Because the IPv6 delivery header does not include a checksum of its own, it is subject to corruption. However, even if the delivery header is corrupted, the likelihood of that corruption resulting in misdelivery of the payload is extremely low.

4.2. MTU Considerations

"IPv6 requires that every link in the Internet have an MTU of 1280 octets or greater. On any link that cannot convey a 1280-octet packet in one piece, link-specific fragmentation and reassembly must be provided at a layer below IPv6" [[RFC2460](#)].

IP adjacencies formed by GRE over IPv6 share this requirement. The IP adjacency MUST have an MTU of 1280 octets or greater. This requirement is fulfilled if all permissible paths between the GRE ingress and GRE egress have PMTU greater than the 1280 plus the GRE overhead.

In case all permissible routes between the GRE ingress and GRE egress do not have PMTU greater than 1280 plus the GRE overhead,

implementations MUST be capable of fragmenting and reassembling the GRE delivery header, as described in [Section 3.1](#).

5. IANA Considerations

This document makes no request of IANA.

6. Security Considerations

This document adds no additional security risks to GRE, beyond what is specified in [[RFC2784](#)]. It also does not provide any additional security for GRE.

7. Acknowledgements

The authors would like to thank Fred Baker, Dino Farinacci, Tom Herbert, Fred Templin, Joe Touch, Andrew Yourtchenko and Lucy Yong for their thorough review and useful comments.

8. Normative References

- [ETYPES] IANA, "ETHER TYPES", 2014,
<<http://www.iana.org/assignments/ieee-802-numbers/ieee-802-numbers.xhtml#ieee-802-numbers-1>>.
- [RFC0791] Postel, J., "Internet Protocol", STD 5, [RFC 791](#), September 1981.
- [RFC1981] McCann, J., Deering, S., and J. Mogul, "Path MTU Discovery for IP version 6", [RFC 1981](#), August 1996.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", [RFC 2460](#), December 1998.
- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", [RFC 2784](#), March 2000.
- [RFC2890] Dommety, G., "Key and Sequence Number Extensions to GRE", [RFC 2890](#), September 2000.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 4443](#), March 2006.

Authors' Addresses

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

Email: cpignata@cisco.com

Ron Bonica
Juniper Networks
2251 Corporate Park Drive
Herndon, Virginia
USA

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

