

Network Working Group
Request for Comments: 1702
Category: Informational

S. Hanks
NetSmiths, Ltd.
T. Li
D. Farinacci
P. Traina
cisco Systems
October 1994

Generic Routing Encapsulation over IPv4 networks

Status of this Memo

This memo provides information for the Internet community. This memo does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Introduction

In an earlier memo [RFC 1701], we described GRE, a mechanism for encapsulating arbitrary packets within an arbitrary transport protocol. This is a companion memo which describes the use of GRE with IP. This memo addresses the case of using IP as the delivery protocol or the payload protocol and the special case of IP as both the delivery and payload. This memo also describes using IP addresses and autonomous system numbers as part of a GRE source route.

IP as a delivery protocol

GRE packets which are encapsulated within IP will use IP protocol type 47.

IP as a payload protocol

IP packets will be encapsulated with a Protocol Type field of 0x800.

For the Address Family value of 0x800, the Routing Information field will consist of a list of IP addresses and indicates an IP source route. The first octet of the Routing Information field constitute a 8 bit integer offset from the start of the Source Route Entry (SRE), called the SRE Offset. The SRE Offset indicates the first octet of the next IP address. The SRE Length field consists of the total length of the IP Address List in octets.

This has the form:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Address Family           | SRE Offset   | SRE Length   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     IP Address List ...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

For the Address Family value of 0xffffe, the Routing Information field will consist of a list of Autonomous System numbers and indicates an AS source route. The third octet of the Routing Information field contains an 8 bit unsigned integer offset from the start of the Source Route Entry (SRE), called the SRE Offset. The SRE Offset indicates the first octet of the next AS number. The SRE Length field consists of the total length of the AS Number list in octets.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Address Family           | SRE Offset   | SRE Length   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     AS Number List ...
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

IP as both delivery and payload protocol

When IP is encapsulated in IP, the TTL, TOS, and IP security options MAY be copied from the payload packet into the same fields in the delivery packet. The payload packet's TTL MUST be decremented when the packet is decapsulated to insure that no packet lives forever.

IP source routes

When a system is processing a SRE with an Address Family indicating an IP source route, it MUST use the SRE Offset to determine the next destination IP address. If the next IP destination is this system, the SRE Offset field should be increased by four (the size of an IP address). If the SRE Offset is equal to the SRE Length in this SRE, then the Offset field in the GRE header should be adjusted to point to the next SRE (if any). This should be repeated until the next IP destination is not this system or until the entire SRE has been processed.

If the source route is incomplete, then the Strict Source Route bit is checked. If the source route is a strict source route and the next IP destination is NOT an adjacent system, the packet MUST be

dropped. Otherwise, the system should use the IP address indicated by the Offset field to replace the destination address in the delivery header and forward the packet.

Autonomous system source routes

When a system is processing a SRE with an Address Family indicating an AS source route, it MUST use the SRE Offset field to determine the next autonomous system. If the next autonomous system is the local autonomous system, the SRE Offset field should be increased by two (the size of an autonomous system number). If the SRE Offset is equal to the SRE Length in this SRE, then the Offset field in the GRE header should be adjusted to point to the next SRE (if any). This should be repeated until the next autonomous system number is not equal to the local autonomous system number or until the entire SRE has been processed.

If the source route is incomplete, then the Strict Source Route bit is checked. If the source route is a strict source route and the next autonomous system is NOT an adjacent autonomous system, the packet should be dropped. Otherwise, the system should use the autonomous system number indicated by the SRE Offset field to replace the destination address in the delivery header and forward the packet. The exact mechanism for determining the next delivery destination address given the AS number is outside of the scope of this document.

Security Considerations

Security issues are not discussed in this memo.

Authors' Addresses

Stan Hanks
NetSmiths, Ltd.
2025 Lincoln Highway
Edison, NJ 08817

EMail: stan@netsmiths.com

Tony Li
cisco Systems, Inc.
1525 O'Brien Drive
Menlo Park, CA 94025

EMail: tli@cisco.com

Dino Farinacci
cisco Systems, Inc.
1525 O'Brien Drive
Menlo Park, CA 94025

EMail: dino@cisco.com

Paul Traina
cisco Systems, Inc.
1525 O'Brien Drive
Menlo Park, CA 94025

EMail: pst@cisco.com

References

RFC 1701
Hanks, S., Li, T, Farinacci, D., and P. Traina, "Generic Routing Encapsulation", RFC 1701, NetSmiths, Ltd., and cisco Systems, October 1994.

