Internet Engineering Task Force
INTERNET-DRAFT
<draft-petri-mobileip-pipe-00.txt>
Date: Jan. 20, 2000

Bernhard Petri
Siemens AG

Expires: July 2000

Private IP Encapsulation within IP (PIPE)

Status of this Memo

Abstract

RFC 2003 specifies a method by which an IP datagram may be encapsulated (carried as payload) within an IP datagram. This is a means to alter the normal IP routing for datagrams, by delivering them to an intermediate destination that would otherwise not be selected by the (network part of the) IP Destination Address field in the original IP header.

This draft allows to extend this encapsulation mechanism also for private IP addresses.

1. Introduction

[RFC 2003] specifies a method by which an IP datagram may be encapsulated (carried as payload) within an IP datagram. Encapsulation is suggested as a means to alter the normal IP routing for datagrams, by delivering them to an intermediate destination that would otherwise not be selected based on the (network part of the) IP Destination Address field in the original IP header.

Once the encapsulated datagram arrives at this intermediate destination node, it is decapsulated, yielding the original IP datagram, which is then delivered to the destination indicated by the original Destination Address field.

[RFC 2003] only allows to encapsulate public IP addresses within IP. However, current IP solutions often use non-unique private IP addresses, e.g. taken from the address space reserved for this purpose in [RFC 1918]. These private IP addresses typically are unique within a private Intranet, e.g. behind a firewall / Network Address Translator (NAT), but are not globally unique and not globally routable in the Internet.

This draft therefore outlines an extension of [RFC 2003] which allows to encapsulate and decapsulate such private IP addresses in the same way as described in [RFC 2003], and to transfer them across the public Internet (also referred to as "tunneling" in RFC 2003). Behaviour not explicitly mentioned in this draft applies as specified in [RFC 2003].

2. Motivation and Solution Overview

2.1 Motivation: Mobility Applications Using Private IP Addresses

It is expected that there may be various applications which will be able to benefit from the PIPE encapsulation mechanism outlined in this draft. An important initial application will be the support of mobile nodes having obtained private IP addresses within a foreign network and/or using private IP addresses in their home network (see also section 2 of [RFC 2003]).

Figure 1 shows an example of a basic tunneling/encapsulation case where a private IP address is translated and encapsulated through a public IP tunnel; within the framework of mobility applications, the source might e.g. be a foreign agent registering at the home agent as the destination, or the source might e.g. be the home agent forwarding data packets to the foreign agent as the destination.

The terms "encapsulator" and "decapsulator" are used as defined in [RFC 2003]. It should be noted that the encapsulator - as the entry point of the tunnel - also performs an address resolution function, in this case between private und public IP addressing. The particular resolution functions used by the encapsulator are outside the scope of this draft.

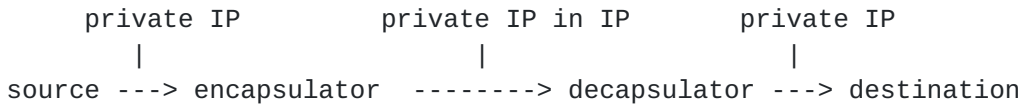

Figure 1: Example for encapsulation of private IP in IP

In Figure 1, the "private IP" address realm between source and encapsulator may or may not be the same "private IP" address realm as the one between decapsulator and destination.

2.2 Identification of Private Addressing Realms

A major problem of the use of private IP addresses is that they are not globally unique, and that a decapsulator receiving an encapsulated private IP address would usually not know, to which address space a private IP address belongs to. In Figure 1 above, the decapsulator will usually not be able to derive the particular private addressing realm from its knowledge about the begin of the tunnel.

Similar problems had already emerged within the Internetwork over NBMA (ION) area of the IETF, and a solution had been developed using a global identification scheme for IP-VPNs which is outlined in [RFC 2685]. This scheme uses an identifier ("VPN-ID") to identify a private network, typically with the objective to provide a related VPN service. The scheme is based on a well-known OUI/index mechanism as e.g. also used for MAC addresses, or for the Interface ID of IPv6 addresses. The VPN-ID consists of a 7-octet format which is split into a 3-octet OUI of a "VPN authority", followed by a 4-octet index allocated by that authority.

This draft proposes to also use this identification scheme for the indication of the particular address realm a private IP address belongs to, when being encapsulated and being transferred across the internet. Since private IP addresses are assumed to be unique within that particular private network, it is easy to attach a VPN-ID to them, e.g. by using the VPN-related OUI of the owner of that network.

2.3. Example: Private IP Interconnection

This example is intended to illustrate and motivate the generic solution specified in sections 3 ff below. In this particular example, the configuration shown in Figure 1 is taken, and it is assumed that both the source and destination belong to the same private address realm "PR1".

Two different cases for the IP-IP tunnels can be distinguished:

a) The decapsulator may be configured in a way that all inner IP header addresses received via IP-IP tunnels, are assumed to belong to one particular IP address space (either a private IP address space or the public IP addressing). In this case, the encapsulator will only insert IP datagrams from that particular address space into the IP-IP tunnel.

b) No particular IP address space is pre-associated with the IP-IP tunnel. This configuration option applies in cases where the encapsulator might send IP datagrams for different address spaces (e.g. public and private) via the IP-IP tunnel to a decapsulator.

In case a), IP-IP encapsulation is applied in a similar way as specified in section 3 of [RFC 2003], with the difference that in this example, the IP addresses within the inner IP header are configured to belong to the private IP address space "PR1" rather than to public IP addressing.

In order to encapsulate a private IP datagram into an IP-IP tunnel in case b), an outer IP header is inserted by the encapsulator before the datagram's existing private IP header, as follows:

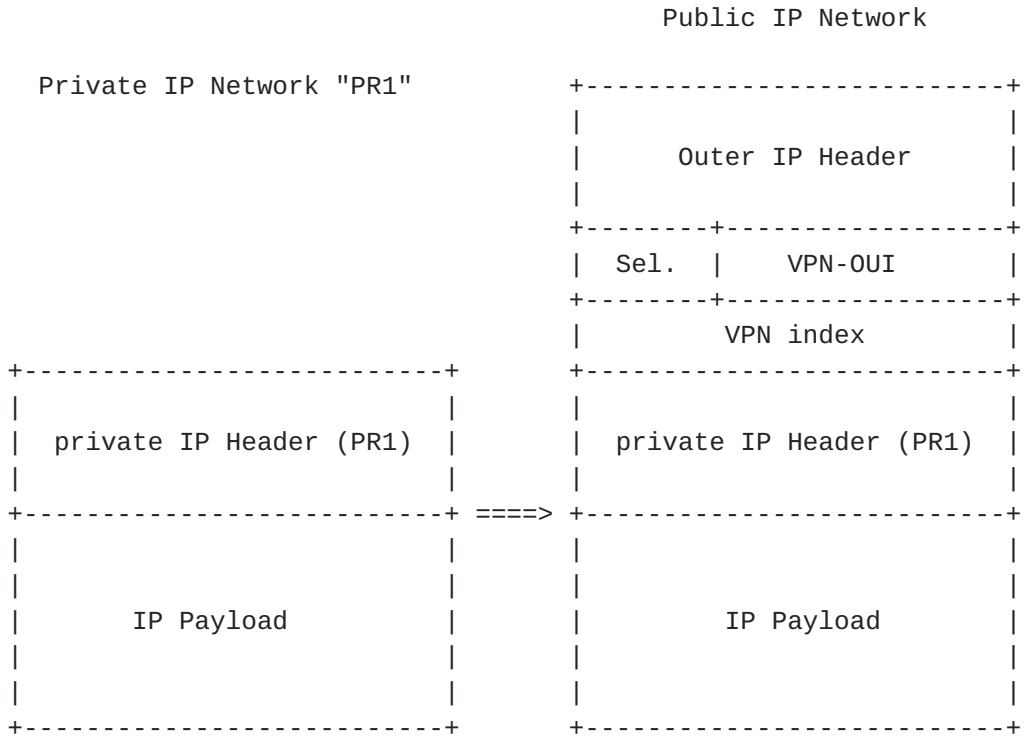

Figure 2: Example for Private IP-IP Encapsulation

The detailed formats for the fields shown in Figure 2 are specified in section 3 below. The Selector field ("Sel.") serves as discriminator / selector indicating which types of addresses are used in the inner IP header. The VPN-Identifier (VPN-ID), consisting of VPN-OUI and VPN index, in this example shows a value identifying PR1.

3. Private IP-IP Encapsulation (PIPE)

3.1 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 [RFC 2119].

In this document, the following definitions and acronyms are used:

Default Addressing Realm -- this is the configured addressing realm a particular (sent or received) source / destination IP address is assumed to belong to, unless another addressing realm is explicitly indicated by a VPN-ID with the mechanisms specified below. In case of [RFC 2003], the default addressing realm is the public IP addressing.

Explicitly Indicated Addressing Realm -- this is the addressing realm specified by a VPN-ID related to a particular sent or received source / destination IP address.

VPN-Identifier (VPN-ID) -- this term is specified in [RFC 2685], and is used in this document for the identification of a particular Explicitly Indicated Addressing Realm.

Other terms, like e.g. "encapsulator", "decapsulator", "inner IP header" and "outer IP header" are used as specified in section 3 of [RFC 2003].

3.2 Configuration Options and Backward Compatibility

This draft outlines a backward compatible extension of [RFC 2003] for the use of private IP addresses. Behaviour not explicitly mentioned in this draft applies as specified in [RFC 2003].

As indicated in [RFC 2003], IP-IP tunnels require knowledge about the decapsulation capabilities at the endpoint of the tunnel. For this document, 3 different configuration options are distinguished which determine how source and destination IP addresses, indicated in the inner IP header, are interpreted.

(1)All source and destination addresses, both in the inner and outer IP headers, are interpreted as being public IP addresses. Support of this configuration option IS REQUIRED for the case of communication with legacy [RFC 2003] devices.

(2)All indicated source and destination IP addresses, indicated in the inner or outer header of IP-IP tunnels, are assumed to belong to one particular Default Addressing Realm. This configuration option MUST e.g. be selected if either the encapsulator or the decapsulator is not able to process VPN-IDs. The Default Addressing Realm may be the public Internet addressing (as in RFC 2003) or any other private IP addressing realm. The Default Addressing Realm for the outer IP header is that of the network between Encapsulator and Decapsulator.

(3) Indicated source IP addresses may belong to an Explicitly Indicated Addressing Realm or to the Default Addressing Realm; indicated destination IP addresses may also belong to an Explicitly Indicated Addressing Realm or to the Default Addressing Realm. This is the most general case.

For equipment complying to this specification, it IS REQUIRED to support configuration options (1) and (2), and it is RECOMMENDED to also support option (3). Support of option (3) IS REQUIRED if more

than one addressing realm is to be supported by IP-IP tunnels.

The main purpose of configuration option (1) is to ensure interoperability with legacy [RFC 2003] devices. For configuration option (1), the IP in IP encapsulation formats specified in [RFC 2003] MUST be used.

The formats for configuration options (2) and (3) are outlined in section 3.3 below. For configuration options (2) and (3), Default Addressing Realms for the inner source and destination IP addresses MUST be preallocated. Note that the Default Addressing Realms for source and destination MAY be different.

For all 3 options, the format of the inner "IP header" and "IP payload" fields MUST be coded as specified in section 3 of [RFC 2003].

3.3 PIPE Encapsulation Formats

For the formats used in configuration options (2) and (3), 5 different cases can be distinguished (see subsections 3.3.1 - 3.3.5 below). For all 5 cases, the "Protocol" field of the outer IP header (identifying the next level protocol) MUST be set to the value: .. <to be allocated by IANA, e.g. value 129>.

The Selector Byte ("Sel."), specified below, is used to distinguish between these cases:

3.3.1 Default Source and Destination Address

If for both source and destination address the Default Addressing Realm is used, the encapsulation format of [RFC 2003] applies. Note that unlike in [RFC 2003], the inner header IP addresses in this case are not automatically assumed to be public IP addresses, but are interpreted according to the Default Address Realm(s). It is not precluded that the Default Address Realms for the source and destination addresses are different.

For configuration option (2), this is the only possible format.

3.3.2 Explicitly Indicated Source and Destination Adress (Same Realm)

If both the source and destination IP address are to be indicated explicitly, and belong to the same addressing realm, the following encapsulation format SHOULD be used:

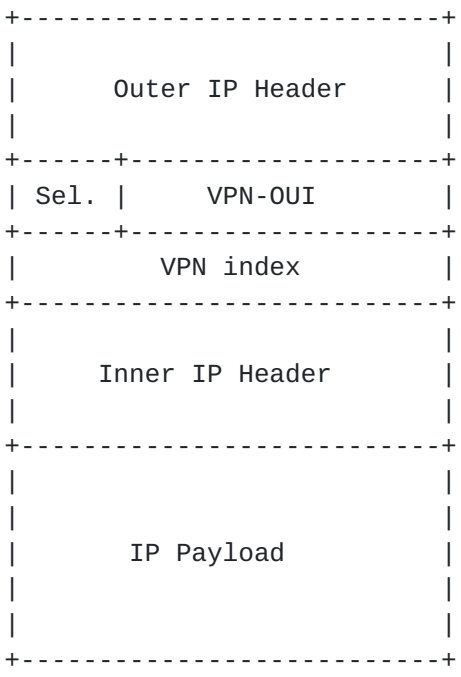

Figure 3: Private IP-IP Encapsulation Format (Case 2)

Outer IP Header: as in [RFC 2003]
Sel.: Selector, value to be allocated by IANA:
tentatively: 0xE0
(= explicitly indicated source / destination address,
same addressing realm)
VPN-OUI, VPN index: as in [RFC 2685]
(refers to both source and destination IP address)

In addition to the format specified above, the format specified in section 3.3.5 below MAY instead by used in certain cases.

3.3.3 Default Source Adress, Explicit Destination Address

If only the destination address is to be indicated explicitly, the following encapsulation format MUST be used:

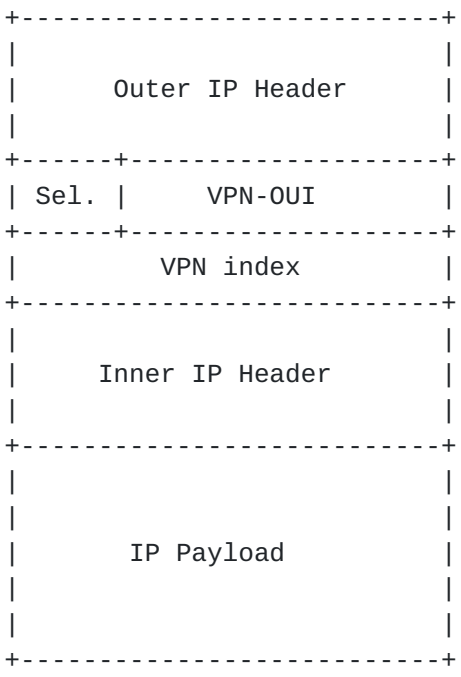

Figure 4: Private IP-IP Encapsulation Format (Case 3)

Outer IP Header: as in [RFC 2003]
Sel.: Selector, value to be allocated by IANA:
tentatively: 0xE1
(=default source, explicitly indicated destination)
VPN-OUI, VPN index: as in [RFC 2685], refers to destination only

3.3.4 Explicit Source Adress, Default Destination Address

If only the source address is to be indicated explicitly, the following encapsulation format MUST be used:

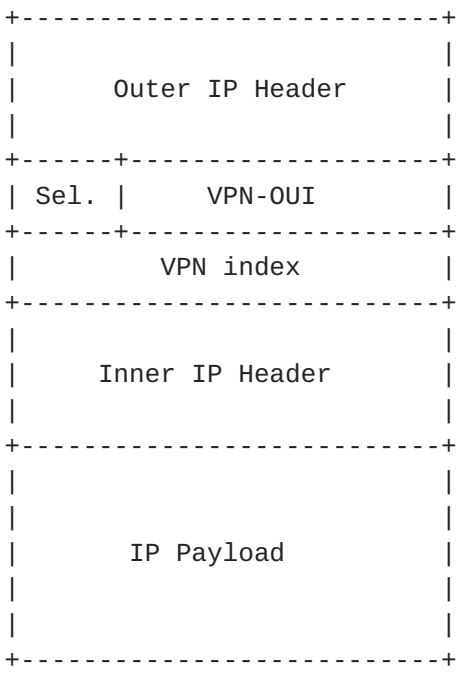

Figure 5: Private IP-IP Encapsulation Format (Case 4)

Outer IP Header: as in [RFC 2003]
Sel.: Selector, value to be allocated by IANA:
tentatively: 0xE2
(=explicitly indicated source, default destination)
VPN-OUI, VPN index: as in [RFC 2685], refers to source only

3.3.5 Explicitly Indicated Source and Destination Address (Different Realms)

If both the source and the destination address are to be indicated explicitly, and belong to different addressing realms, the following encapsulation format MUST be used:

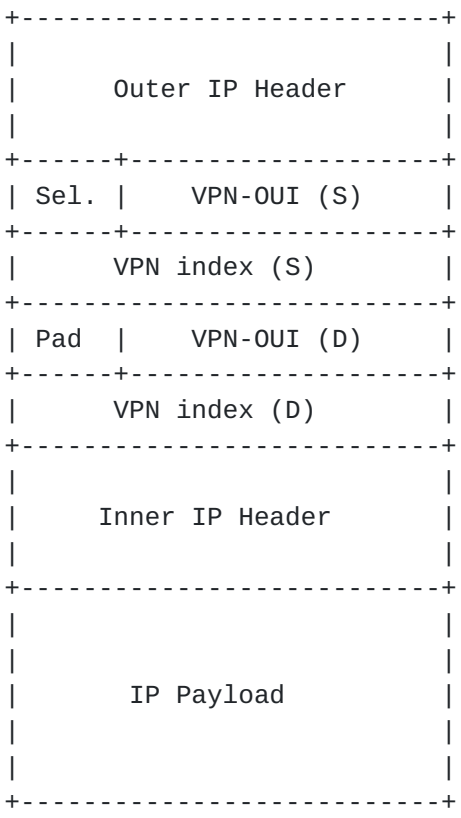

Figure 6: Private IP-IP Encapsulation Format (Case 5)

Outer IP Header: as in [RFC 2003]
Sel.: Selector, value to be allocated by IANA:
tentatively: 0xE3
(=explicitly indicated source and destination,
different addressing realm)
VPN-OUI (S), VPN index (S): as in [RFC 2685], refers to source
Pad: Pad field (inserted for 32-bit alignment), this field
MUST be coded as 0x00, and is ignored on receipt)
VPN-OUI (D), VPN index (D): as in [RFC 2685], refers to
destination

4. Example: A Mobile Node Registers at its Home Agent

This section provides an example illustrating how the formats specified in section 3 above can be applied. In this example, a mobile node MN has obtained a temporary private IP address within a private IP address realm PR2, where it is currently located. It now wants to register this address with its home agent HA which owns a private IP address within realm PR3.

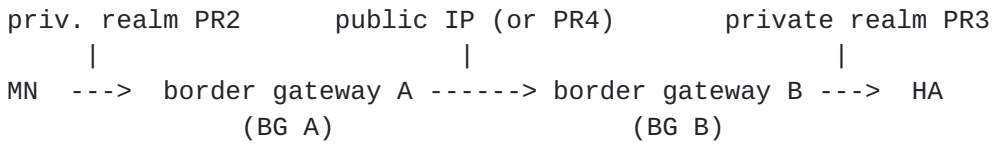

Figure 7: Example for encapsulation of private IP in IP

Since PR3 is different from PR2, the following Private IP in IP Encapsulation (PIPE) formats are used on the way from MN to HA:

Step 1: From MN to BG A:

Outer Header:	Source Address:	MN(PR2)
	Destination Address:	BG A(PR2)
Inner Header:	Source Address:	MN(PR2) = default
	Destination Address:	HA (PR3)
Selector:	0xE1 (explicitly indicated destination)	

Step 2: From BG A to BG B:

Outer Header:	Source Address:	BG A(public IP or PR4)
	Destination Address:	BG B(public IP or PR4)
Inner Header:	Source Address:	MN (PR2)
	Destination Address:	HA (PR3)
Selector:	0xE3 (explicit source and destination, different address realms)	

Step 3: From BG B to HA:

Outer Header:	Source Address:	BG B(PR3)
	Destination Address:	HA(PR3)
Inner Header:	Source Address:	MN(PR2)
	Destination Address:	HA(PR3) = default
Selector:	0xE2 (explicitly indicated source)	

It should be noted that in addition to the formats illustrated above, a real transfer of a packet from MN to HA also involves a number of routing decisions, and address resolution functions which are outside the scope of this specification, and may perhaps be specified in separate drafts.

In the example above, the use of the public Internet as a backbone to interconnect the address realms PR2 and PR3 seems to be a natural

choice, but another private realm (PR4) might also be suitable for that purpose. The selection of a transit backbone or a particular address resolution may be subject to different criteria, and/or may be dependent on particular applications.

5. Security Considerations

IP encapsulation potentially reduces the security of the Internet, and care needs to be taken in the implementation and deployment of IP encapsulation. More detailed considerations of security implications of IP-IP tunnels can be found in section 6 of [RFC 2003].

Since private addresses are typically administered to prevent access to networks inside an enterprise, the transfer of private addresses across networks outside this enterprise must be handled with great care. It may be required to use authentication and possibly encryption to maintain the existing security policy which originally dictated the choice of using a private address space within the enterprise.

6. IANA Considerations

It is proposed that IANA establishes and maintains a list of protocol values for the selector byte following the outer IP header. The following values are an example for this possible list:

0x00 through 0xDF	as for corresponding IP version number and header field
0xE0 - 0xE3	as defined in this specification
0xE4 - 0xFF	reserved

In addition, the formats for private IP-IP encapsulation specified in this document require the allocation of a new value in the "Protocol" field (identifying the next header) within the IPv4 header, (e.g.: 129 PIPE Private IP-IP Encapsulation ....).

7. IPR Considerations

Siemens may own intellectual property on some of the technologies described in this document.

References

[RFC 1918] Rekhter, Y. et al., "Address Allocation for Private Internets", RFC 1918, Febr. 1996

[RFC 2003] Perkins, C., "IP Encapsulation within IP", RFC 2003, October 1996

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

[RFC 2685] Fox, B., Gleeson, B., "Virtual Private Networks Identifier", RFC 2685, September 1999.

Author Information

Bernhard Petri
Siemens AG
Hofmannstr. 51
Munich, Germany, D-81359
phone: +49 89 722-34578
email: bernhard.petri@icn.siemens.de