

Source Demand Routing Working Group  
INTERNET DRAFT

Peter Ford  
LANL  
Tony Li  
cisco Systems  
Yakov Rekhter  
T.J. Watson Research Center, IBM Corp.  
October 1994

**Explicit Routing Protocol (ERP) for IPv6**  
**<[draft-ietf-sdr-erp-00.txt](#)>**

Status of this Memo

This document is an Internet Draft. Internet Drafts are working documents of the Internet Engineering Task Force (IETF), its Areas, and its Working Groups. Note that other groups may also distribute working documents as Internet Drafts.

Internet Drafts are draft documents valid for a maximum of six months. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material or to cite them other than as a ``working draft'' or ``work in progress.''

Please check the `1id-abstracts.txt` listing contained in the internet-drafts Shadow Directories on `nic.ddn.mil`, `nnsc.nsf.net`, `nic.nordu.net`, `ftp.nisc.sri.com`, or `munari.oz.au` to learn the current status of any Internet Draft.

## **1. Introduction**

This document specifies the format and processing of the IPv6 Explicit Routing Protocol (ERP) Header. The purpose of this header is to provide IPv6 forwarding functionality comparable to SDRP [[SDRP](#)]. The reader should be familiar with [IPv6].

## **2. Acknowledgments**

This document is based on "Source Demand Routing: Packet Format and Forwarding Specification (Version 1)" [[SDRP](#)] and "Source Demand Routing: Route Setup" [[SDRP-SETUP](#)].

## **3. Model of Operation**

Expiration Date April 1995

[Page 1]

An Internet can be viewed as a collection of routing domains interconnected by means of common subnetworks, and Border Routers (BRs) attached to these subnetworks. A routing domain itself may be composed of further subnetworks, routers interconnecting these subnetworks, and hosts.

For the purposes of this document, a BR belongs to only one domain. A pair of BRs, each belonging to a different domain, but attached to a common subnetwork, form an inter-domain connection. By definition, packets that traverse multiple domains must traverse BRs of these domains. Note that a single physical router may act as multiple BRs for the purposes of this model.

A connected sets of Routing Domains may be grouped into Routing Domain Confederations. A domain may belong to more than one confederation. The confederations may be nested, disjoint, or overlapped.

### **3.1 Addressing Assumptions**

Each domain has a globally unique identifier, called a Routing Domain Identifier (RDI). All the BRs within a domain need to know the RDI assigned to the domain. Likewise, each routing domain confederation has a globally unique identifier, called a Routing Domain Confederation Identifier (RDCI).

This document assumes that RDIs and RDCIs are expressed as IPv6 addresses, and assigned based on the unicast address prefixes assigned to domains and confederations.

A subscriber RD should use as its RDI an address taken out of the unicast address prefix assigned to it. A multihomed RD should use as its RDI an address taken out of one of the unicast address prefixes assigned to it. A service provider should use as its RDI an address taken out of the unicast address prefix that the provider uses for assigning addresses to systems within the provider. Finally, an RDI of a domain may be constructed by taking a unicast address out of any unicast address prefix assigned to the domain.

If a service provider forms a Routing Domain Confederation with some of its subscribers and the subscribers take their addresses out of the provider, then an address taken out of the unicast address prefix assigned to the provider should be used as the RDCI of the confederation. An RDCI of a confederation may be also constructed by taking a unicast address out of any unicast address prefix assigned to any domain within the confederation, provided that this

unicast address is different from any other RDI or RDCI.

The address taken out for RDI or RDCI can not be used for unicast address assignment.

Expiration Date April 1995

[Page 2]

The assumptions made here, combined with the assumption that unicast addresses are globally unambiguous, guarantee that non-multicast IPv6 addresses (except for the block allocated for private use) are globally unambiguous and that each assigned address can represent either (a) unicast address, or (b) RDI, or (c) RDCI.

A router may have one or more unicast addresses, one RDI, one or more RDCIs associated with it. It is assumed that at the minimum each router has to have at least one unicast address, but may also have at most one domain identifier and one or more confederation identifiers.

For the rest of the document, when we refer to an address of a router this reference may include either a unicast address or RDI or RDCI associated with the router.

### **3.2 ERP Route Definition**

An ERP route is defined as a sequence of elements (called forwarding hops), where each hop may depict a router, a domain, or a confederation. Each hop is syntactically expressed as a unicast address. Thus, an ERP route is syntactically expressed as a sequence of unicast addresses. Semantically an ERP route is defined as an arbitrary intermix of routers, domains, and confederations.

A hop within an ERP route can either be "strict" or "loose". Forwarding to a strict next hop requires that the next hop be adjacent to the current hop. Forwarding to a loose next hop doesn't impose this requirement.

The definition of what constitute an "adjacent" hop depends on the type (i.e., unicast address, RDI, RDCI) of the current and the next hop, and is defined as follows.

A hop with unicast address A is adjacent to a hop with unicast address B, if there are elements N and M, N with unicast addresses A and X, and M with unicast addresses B and Y, and addresses X and Y share a common subnetwork.

A hop with unicast address A is adjacent to a hop with domain identifier B, if there is an element N with domain identifier B and unicast address X, so that A is adjacent to X (as defined above). In this case any element with domain identifier B is adjacent to A.

A hop with a unicast address A is adjacent to a hop with confederation identifier B if there is an element N with confederation identifier B and unicast address X, so that A is adjacent to X (as defined above).

A hop with domain identifier A is adjacent to a hop with domain identifier B if there exist two elements, M and N, M with domain identifier A and unicast address X, and N with domain identifier B and unicast address Y, such that X is adjacent to Y (as defined

above).

A hop with domain identifier A is adjacent to a hop with confederation identifier B if there exist two elements, M and N, M with domain identifier A and unicast address X, and N with confederation identifier B and unicast address Y, such that X is adjacent to Y (as defined above).

A hop with confederation identifier A is adjacent to a hop with confederation identifier B if there exist two elements, M and N, M with confederation identifier A and unicast address X, and N with confederation identifier B and unicast address Y, such that X is adjacent to Y (as defined above).

Note that the notion of "adjacent" is reflexive. That is, if A is adjacent to B, then B is adjacent to A.

### **3.3 Setup**

In some cases, commonly known as "flows", where the duration of a packet stream is significantly longer than the end-to-end round-trip delay, and particularly where the amount of payload in the packet is small, it may be worthwhile to "set up" the ERP route by saving state information in routers that have to process the route, instead of carrying the full ERP route in every packet. Once this state is established, the originator of the ERP route can use a combination of Source Address and Flow Label carried in the fixed header to refer to the ERP route rather than carrying the full route in each data packet, thereby reducing the ERP Header size and transmission time.

It is important to the protocol that it does not impose setup on routers that are short on state space or that otherwise restrict setup. Therefore, the desire for setup is simply flagged by the the originator of the ERP route, and the routers that have to process the ERP route may choose to accept or reject the request. If all of the routers that have to process the ERP route accept, then the originator can begin using a route identifier (as carried by the Flow Label field of the Fixed Header) to refer to the saved state. If any router refuses setup, the originator must continue including the full ERP route in each data packet or else try a different route.

When a router rejects a setup request, it sends an ICMP message containing the route identifier to the originator of the request. ICMP messages are also used in this manner if a router loses or is missing state for a particular route identifier. When a router accepts a setup request, it continues forwarding the request along

the ERP route. Successful route setup is indicated when the final router on the ERP route sends an ICMP message containing Flow Label to the originator of the route.

### 3.4 Forwarding Information Base (FIB) Model

It is assumed that each router maintains a conventional unicast Forwarding Information Base (FIB), where each entry in the FIB is represented as a tuple <address prefix, next hop>.

It is assumed that if a router shares a common subnetwork with a BR, then the router has information about all the RDIs and RDCIs the BR belongs to. The mechanisms by which this information is acquired are outside the scope of this document. This information is represented in the FIB as a set of tuples, where the address prefix component of each tuple specifies a particular RDI or RDCI (expressed as a 16 octets IPv6 address) of the BR and the next hop component specifies the IPv6 unicast address of the BR. This also applies to the case where the router is a BR itself.

If the next hop depicts a router that is not directly connected to the local system (doesn't share a common subnetwork), then it is assumed that the FIB contains an entry that allows to derive the immediate next hop (the next hop on a common subnetwork).

In addition to conventional FIB a router may maintain a Flow Label cache. This cache is used to support setups. Each entry in the cache is a triplet <Source Address, Flow Label, ERP Header>. An entry in the cache is uniquely identified by its the first two components of the triplet (Source Address and Flow Label). Each entry in the cache can be timed out at any time, and should be timed out periodically. The particular policy used for timing out cache entries is a local matter.

## 4. Explicit Forwarding Header format

An ERP Header is a type of an IPv6 Routing Header. The value of the Routing Type field of the IPv6 Routing Header is set to 1 for an ERP Header. The ERP Header format is defined as following:

```

      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  | Routing Type=1 | Flags   | NestL|   FwdDirLen  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| NextHopPtr   |               Strict/Loose Bit Mask               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

```
|      Forwarding Directions, a list of IPv6 addresses      |
|      (integral multiple of 16 octets)                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
```

Next Header - 8-bit selector.

Identifies the type of the header immediately following the Routing header. Uses the same values as the IPv4 Protocol field [[RFC-1340](#)] plus new types defined for IPv6.

Routing Type - 1.

Flags (5 bits):

MRE, Must Report Errors (bit 0).

If this bit is set to 1, and a router can not further forward a packet, the router must generate an ICMP error message. If this bit is set to 0, and a router can not further forward a packet the router should not generate an ICMP error message.

FFD, Behavior on Failure of Forwarding Directions (bit 1).

If this bit is set to 1, it indicates that if a router can not further forward a packet the router must abort the ERP route (as specified in [Section 7](#)).

SAS, Source Address Swap (bit 2).

If this bit is set to 1, it indicates that the entity that added the current ERP Header to a packet placed the source address from the fixed header as the first element in the Forwarding Directions, and placed its own unicast address as the source address in the fixed header.

''' in 14

FLS, Flow Label Saved (bit 4)

If this bit is set to 1, then the low-order 28 bits of the second element in the Forwarding Directions field contain a Flow Label.

Nesting Level (3 bits).

The Nesting Level field is used to control the number of ERP Headers carried by a packet. The maximum number of ERP Headers that can be carried by a packet is 8.

Forwarding Directions Length (1 octet unsigned integer).

The Forwarding Directions Length field contains the number

of Forwarding Directions hops in the header. Length of the header can be calculated from this value (length = FwdDirLen \* 16 + 8) The value of this field shall not exceed 24.

Next Hop Pointer (1 octet unsigned integer).

The Next Hop Pointer field contains the index of next hop to be processed. The first element in the Forwarding Directions has index 0. The value of this field shall not exceed 24.

Strict/Loose Bit Mask (24 bits).

The Strict/Loose Bit Mask is used when making a forwarding decision. If the N-th bit (N not equal 1) in the Strict/Loose Bit Mask field is set to 1, it indicates that the N-th element in the Forwarding Directions is a Strict Hop (with respect to the (N-1) element). If this bit is set to 0, it indicates that the N-th element in the Forwarding Directions is a Loose Hop (with respect to the (N-1) element). The value of the first bit in the Strict/Loose Bit Mask field specifies whether the first element in the Forwarding Directions is a Strict or a Loose Hop with respect to the entity that originated an ERP route.

Forwarding Directions (variable length, multiple of 16 octets).

The elements of the Forwarding Directions (hops) are syntactically IPv6 addresses. Semantically Forwarding Directions can contain an arbitrary intermix of unicast addresses, RDIs, and RDCIs.

## **5. Originating ERP Routes**

This document assumes that an entity that attaches an ERP Header to a packet is preconfigured with a set of ERP routes. Procedures for constructing these routes are outside the scope of this document. Use of ERP capabilities may be deployed initially without additional routing information exchange protocol support.

An entity that is capable of attaching ERP Headers is expected to use information carried in a packet combined with the local criteria to determine whether the packet should be forwarding along a particular ERP route. Associated with each set of criteria is a set of one or more ERP routes that should be used to route matching packets. The exact nature of the criteria is a local matter.

An ERP Header may be attached to a packet by either an entity that originated the packet or by an entity that forwards the packet. For example, either a host that originates a packet, or any router that

forwards the packet may attached an ERP Header to the packet. An ERP Header is attached to a packet by placing the Header immediately after the fixed header.

A packet may carry more than one ERP Headers. Each such header is

referred to by its nesting level (as carried in the Nesting Level field of the header). The first Header attached to a packet is referred to as the level 1 ERP Header. The level 1 ERP Header is identified by the Next Header field in the ERP Header (the value of the field is other than Routing Header). The outermost ERP Header (the one that immediately follows the fixed header) is referred to as the "current" ERP Header.

When an entity attaches the level 1 ERP Header to a packet, the entity must put the destination address, as specified in the destination address field of the fixed header of the packet, as the last element in the Forwarding Directions of the ERP Header. Thus, the last element in the Forwarding Directions field of the level 1 ERP Header is the destination address of the packet.

Any router may attach a level 1 ERP Header to a packet that the router has to forward. Only a router that is identified by the current element in the Forwarding Directions field of the current ERP Header can attach a non-level 1 ERP Header to the packet, and only when the next element in the Forwarding Directions is specified as a Loose Hop.

The entity that attaches the level 1 ERP Header to a packet specifies the maximum allowed number of nested ERP Headers. When a router wants to attach an ERP Header to a packet that already has one or more ERP Headers the router decrements the Nesting Level from the current ERP Header by 1 and places it in the Nesting Level field of the ERP Header it wants to attach. No further ERP Headers can be attached to a packet whose current ERP Header carries zero as the value of the Nesting Level field.

When an entity other than the one that originates a packet, attaches an ERP Header to the packet, the entity places the source address from the fixed header as the first element in the Forwarding Directions field, places its own unicast address in the source address field of the fixed header, and sets the Source Address Swap flag to 1. If the value of the Flow Label field in the fixed header is non-zero, then this value is placed in the low-order 28 bits of the second element in the Forwarding Directions field, the Flow Label Save flag is set to 1, and the value of the NextHopPtr field is set to 2. Otherwise (the value of the Flow Label field is zero), the value of the NextHopPtr field is set to 1. If the Flow Label is zero, then the actual ERP route starts from the second element of the Forwarding Directions, otherwise the route starts from the third element of the Forwarding Directions.

If an entity that originates a packet attaches an ERP Header to the

packet, then the entity must set the Source Address Swap flag to 0, Flow Label Save flag to 0, and set the value of the NextHopPtr to 0.

We say that a packer carries an implicit ERP Header if the packet's forwarding is done using the setup capabilities, and the actual ERP Header that was used during the setup is omitted. Otherwise, 'we

say that the Header is explicit.

An entity originating an ERP route (the originator) may request an ERP setup by setting the appropriate bits in the Flow Label field, and setting all other fields as described above. It is suggested to set the Must Report Errors flag to 1.

The originator may wait a full round-trip time (provided that this time is known by the originator) for a response to the setup request, in the meantime sending subsequent packets with an explicit ERP Header. There is no limit, however, to the number or frequency of setup requests, thus the entity may repeat the setup. Once the originator receives a response indicating that the setup is successfully completed the originator may start sending packets with an implicit ERP Header. The originator may choose to send packets with the implicit ERP Header immediately after sending the setup request. This may be useful if the packets will be useless after waiting a full round-trip time. In this case, the packets will be delivered if the setup is successful, but may be dropped otherwise.

## **6. Processing packets with ERP routes**

We say that a router receives a packet with an ERP route if the destination address in the fixed header of the packet that arrives at the router depicts one of the addresses of the router and the packet carries at least one ERP Header.

In the case where a router receives a packet with an explicit ERP Header we say that the packet carries a completed ERP route if either (a) the Header is Level 1 and the value of the NextHopPtr field in the Header is one less than the value of the FwdDirLen field, or (b) the Header is other than Level 1 and the value of the NextHopPtr field in the Header is equal to the value of the FwdDirLen field. Otherwise, we say that the received packet carries an incomplete (non-empty) ERP route.

In the case where a router receives a packet with an implicit ERP Header the router checks for a presence of an entry in its Flow Label cache whose Source Address is equal to the Source Address field of the fixed header of the packet, and whose Flow Label is equal to the Flow Label field of the fixed header. If no such entry is found, the packet is processed as described in [Section 7](#). If such an entry is found, then we say that the packet carries a completed ERP route if the ERP Header associated with the entry is either (a) a Level 1 and the value of the NextHopPtr field in the Header is one less than the value of the FwdDirLen field, or (b) is other than Level 1 and the value of the NextHopPtr field in the Header is equal to the value of

the FwdDirLen field. If such an entry is found, but neither (a) nor (b) are satisfied, then we say that the received packet carries an incomplete (empty) ERP route.

### **6.1 Processing packets with a completed ERP route**

When a router receives a packet with an explicit ERP Header and the ERP route is completed then the packet is processed as follows.

If the current ERP Header is a Level 1 header, then the last element in the Forwarding Directions is placed in the Destination Address field of the fixed header. If the Source Address Swap flag in the current ERP Header is set to 1, then the first element in the Forwarding Directions field is swapped with the Source Address field of the fixed header. If the Flow Label Save flag in the current ERP Header is set to 1, then the low-order 28 bits of the second element in the Forwarding Directions field is placed in the Flow Label field of the fixed header. If the current ERP Header is not a Level 1 header, then the current header must be removed from the packet. If the current ERP Header is a Level 1 header, then it may be removed from the packet. In addition, if the Flow Setup Request flag in the Header is set to 1, and the router is willing to accommodate the setup request (as described in [Section 6.4](#)) the router generates an ICMP message to the entity depicted by the source address in the fixed header indicating that the setup request has been successfully completed.

When a router receives a packet with an implicit ERP Header and the ERP route is completed then the packet is processed as described in the previous paragraph, except that instead of the current ERP Header the router uses the ERP Header associated with the entry in the Flow Label cache identified by the Source Address and Flow Label fields of the packet's fixed header.

After the above procedures are completed the packet is submitted for normal forwarding procedure.

### **6.2 Processing packets with an incomplete ERP route and explicit ERP Header**

If a router receives a packet with an incomplete ERP route, and the destination address in the fixed header is one of the unicast addresses associated with the router, the router must check whether the current element in the Forwarding Directions of the current ERP Header, as pointed by the NextHopPtr, specifies an address associated with the router. If this condition isn't satisfied, then further handling of the packet is done as described in [Section 7](#). If this condition is satisfied, then the packet is processed as follows.

If the current ERP Header is a Level 1 header and the value of the

NextHopPtr field is two less than the value of the FwdDirLen field, or if the current ERP Header is not a Level 1 header and the value of the NextHopPtr field is one less than the value of the FwdDirLen field, the router increments the value of the NextHopPtr field by 1, and then follows the procedures specified in [Section 6.1](#).

Otherwise, the router extracts the next element from the Forwarding Directions (the next element is the element that immediately follows the current element in the Forwarding Directions) and the strict/loose indicator associated with the next element (the indicator is extracted from the Strict/Loose Bit Mask field). Further processing depends on whether the next element specifies a Loose Hop (see [Section 6.2.1](#)) or a Strict Hop (see [Section 6.2.2](#)).

#### **[6.2.1](#) Loose Next Hop**

If the next element in the Forwarding Directions is a Loose Hop, then the router places the address, as specified in the next element, in the destination address field of the fixed header, increments the value of the Next Hop Pointer by 1, and submits the packet for normal forwarding procedure. If the forwarding procedure can not forward the packet (based on the destination address specified in the fixed header), then further handling of the packet is done as described in [Section 7](#).

#### **[6.2.2](#) Strict Next Hop**

If the next element in the Forwarding Directions is a Strict Hop, then the router performs a lookup in its FIB (using the "longest match" algorithm) using the address specified in the next element as an index. If no matching tuple is found, then further handling of the packet is done as described in [Section 7](#).

If the next hop component of the found tuple specifies an entity that shares a common subnetwork with the router, then the router checks whether any of the addresses associated with that entity are equal to the address specified by the next element in the Forwarding Directions. If that is the case, then the value of the Next Hop Pointer field is incremented by 1.

The address specified by the next hop component of the found tuple is placed in the destination address field of the fixed header and the packet is then submitted for normal forwarding procedure. If the forwarding procedure can not forward the packet (based on the destination address specified in the fixed header), then further handling of the packet is done as described in [Section 7](#).

### **[6.3](#) Processing packets with an incomplete ERP route and implicit ERP**

Header

If a router receives a packet with an implicit ERP Header that carry an incomplete ERP route, and the destination address in the fixed

Expiration Date April 1995

[Page 11]

header is one of the unicast addresses associated with the router, the router extracts from its Flow Label cache an entry identified by the Source Address and Flow Label fields of the packet's fixed header.

If the ERP Header associated with the extracted entry is a Level 1 header and the value of the NextHopPtr field is two less than the value of the FwdDirLen field, or if the Header is not a Level 1 header and the value of the NextHopPtr field is one less than the value of the FwdDirLen field, the router follows the procedures specified in [Section 6.1](#).

Otherwise, the router extracts the next element from the Forwarding Directions (the next element is the element that immediately follows the current element in the Forwarding Directions) and the strict/loose indicator associated with the next element (the indicator is extracted from the Strict/Loose Bit Mask field). Further processing depends on whether the next element specifies a Loose Hop (see [Section 6.3.1](#)) or a Strict Hop (see [Section 6.3.2](#)).

#### [6.3.1](#) Loose Next Hop

If the next element in the Forwarding Directions is a Loose Hop, then the router places the address, as specified in the next element, in the destination address field of the fixed header, and submits the packet for normal forwarding procedure. If the forwarding procedure can not forward the packet (based on the destination address specified in the fixed header), then further handling of the packet is done as described in [Section 7](#).

#### [6.3.2](#) Strict Next Hop

If the next element in the Forwarding Directions is a Strict Hop, then the router performs a lookup in its FIB (using the "longest match" algorithm) using the address specified in the next element as an index. If no matching tuple is found, then further handling of the packet is done as described in [Section 7](#).

If the next hop component of the found tuple specifies an entity that shares a common subnetwork with the router, then the router checks whether any of the addresses associated with that entity are equal to the address specified by the next element in the Forwarding Directions.

The address specified by the next hop component of the found tuple is

placed in the destination address field of the fixed header and the packet is then submitted for normal forwarding procedure. If the forwarding procedure can not forward the packet (based on the destination address specified in the fixed header), then further handling of the packet is done as described in [Section 7](#).

#### **6.4 Processing packets with Setup Request**

If a router receives a packet with an ERP route and the packet carries a setup request, then in addition to the procedures described in Sections [6.1](#) and [6.2](#) (including 6.2.1 and 6.2.2), the router performs the following procedures.

If the router is willing to accommodate the setup request, then the router constructs a triplet <Source Address, Flow Label, ERP Header>, where the value of the first element (Source Address) is set to the value of the Source Address from the fixed header, the value of the second element (Flow Label) is set to the value of the Flow Label field of the fixed header, and the value of the third element is set to the value of all the ERP Headers carried by the packet. The router then checks whether the triplet is present in its Flow Label cache, and if not then adds to the cache an entry that contains the triplet.

If the router is unwilling to accommodate the setup request, and Must Report Error flag is set to 1, then the router generates an ICMP message to the entity depicted by the source address in the fixed header indicating that the setup request can not be accommodated.

The rules by which a router decides whether to accommodate a setup request are outside the scope of this document.

### **7 Error Handling**

If the MRE is set to 1, then the router generates an ICMP error message back to the entity specified in the source address of the fixed header. If the MRE is set to 0, then no ICMP error messages should be generated.

If the Behavior on Failure of Forwarding Directions bit in the current ERP Header of a packet is set to 1, and the Header carries a non-empty ERP route then the packet is processed as follows. If the current ERP Header is a Level 1 header, then the value of the NextHopPtr field is set to one less than the value of the FwdDirLen field. Otherwise, the value of the NextHopPtr field is set to the value of the FwdDirLen field. The packet is then processed as described in [Section 6.1](#).

If the Behavior on Failure of Forwarding Directions bit in the current ERP Header of a packet is set to 1, the Header carries an empty ERP route, and the Header is not a level 1 Header, then the packet is processed as described in [Section 6.1](#). If the Behavior on

Failure of Forwarding Directions bit in the current ERP Header of a packet is set to 1, the Header carries an empty ERP route, and the Header is a level 1 Header, then the packet is discarded.

If the Behavior on Failure of Forwarding Directions bit in the

current ERP Header of a packet is set to 0, then the packet is discarded.

## **8 Path MTU Handling**

## **9. ERP ICMP Message Format**

[Need to define format of an ICMP message associated with ERP. The message should be able to carry one of the following indications:

- (a) Setup Request completed
- (b) Setup Request refused
- (c) Can't satisfy Forwarding Directions
- (d) No Flow Label cache entry found

In addition to the indication the message should carry as much of the original packet as possible, but at the minimum the fixed header and all the ERP Headers.]

## **10. Authors' Address**

Peter Ford  
Los Alamos National Laboratory  
Los Alamos  
Phone: (505) 665-0058  
e-mail: peter@goshawk.lanl.gov

Tony Li  
cisco Systems, Inc.  
1525 O'Brien Drive  
Menlo Park, CA 94025  
email: tli@cisco.com

Yakov Rekhter  
T.J. Watson Research Center IBM Corporation  
P.O. Box 704  
Yorktown Heights, NY 10598  
Phone: (914) 784-7361  
email: yakov@watson.ibm.com

## **7. References**

Expiration Date April 1995

[Page 14]

[RFC-1340] Reynolds, J. and Postel, J. "Assigned Numbers". July 1992.  
[RFC 1340](#).

[SDRP] Estrin, D., Li, T. , Rekhter, Y., and Varadhan K., "Source Demand Routing: Packet Format and Forwarding Specification (Version 1)" 22 March, 1994. Internet Draft.

[SDRP-SETUP] Estrin, D., Zappala, D., Li, T., "Source Demand Routing: Route Setup", June 1993, Internet Draft.

[SIPP-16] Deering, S., "Simple Internet Protocol Plus (SIPP) Specification (128 bit address version)". 17 July 1994. Internet Draft.

Expiration Date April 1995

[Page 15]