Network Working Group
Internet-Draft
Expires: April 29, 2003

M. Bagnulo
A. Garcia-Martinez
UC3M
October 29, 2002

Extension Header for Site-Multi-homing support
draft-bagnulo-multi6-mhexthdr-00

Status of this Memo

This Internet-Draft will expire on April 29, 2003.

Copyright Notice

Abstract

This note describes an IPv6 multi-homing solution that achieves equivalent fault tolerance benefits to those provided by current IPv4 multi-homing solution while preserving the route aggregation capabilities of the Provider-based Aggregation scheme. The solution lies on the inclusion, in every packet flowing to a multi-homed site, of an extension header containing multiple alternative route information to the destination, so that if the original destination address becomes unreachable, alternative route is used for reaching the destination.

1. Introduction

This note describes an IPv6 multi-homing solution that achieves equivalent fault tolerance benefits to those provided by current IPv4 multi-homing solution while preserving the route aggregation capabilities of the Provider-based Aggregation scheme. The solution lies on the inclusion, in every packet flowing to a multi-homed site, of an extension header containing multiple alternative route information to the destination, so that if the original destination address becomes unreachable, alternative route is used for reaching the destination. Additionally, a Destination option is defined (the Alternative Prefix Destination Option) to convey multiple alternative prefix information from a multi-homed host to the other end of the communication.

2. Solution components

2.1 Alternative Prefix Destination option

The following destination option is defined:

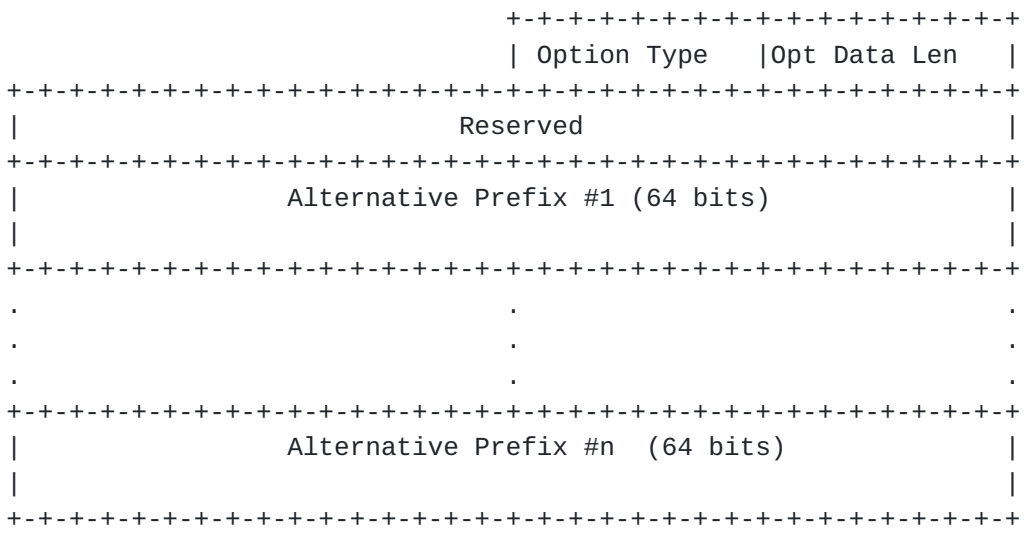

Option Type value is 000xxxxx (in bits): The two highest order bits set to O, so that if the option is not recognized, the option is ignored and the packet is processed [1]. This allows that hosts not implementing this solution to be capable of communication with hosts which do implement the solution. Note that multi-homing benefits are lost in this particular communication. The third bit is set to 0, since the option data does not change in the route [1]. Remaining bits are set to xxxxx (TBD by IANA)

Option Data Length value is 4n+2, since the option contains n Alternative Prefixes, and each one has 4 octets and 2 octets to preserve alignment.

Alternative Prefix field contains alternative Prefixes assigned to the source interface other than the one included in the Source Address field of the IPv6 header [1].

The intended use of the above destination option is the communication of multiple alternative routes to the multi-homed site, from a multi-homed source node to a destination node.

2.2 Alternative Prefix Extension Header

A new Extension Header is defined with the following format:

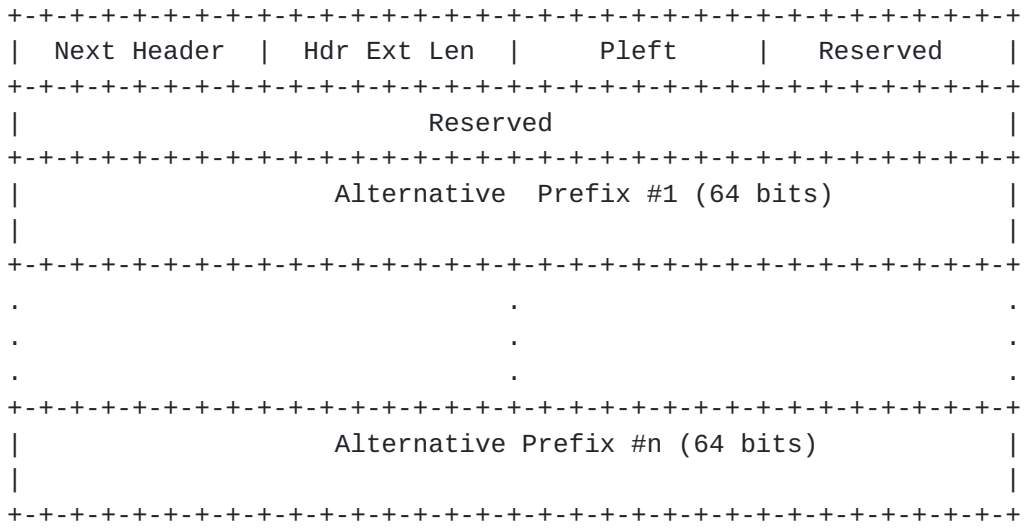

Next header value: 8-bit selector. Identifies the type of header immediately following the Alternative Prefix Extension Header.

Hdr Ext Len: 8-bit unsigned integer. Length of the Alternative Prefix Extension Header in 8-octet units, not including the first 8 octets.

Pleft: 8-bit unsigned integer. Number of Alternative Prefixes left, i.e., number of Prefixes that have not been used for reaching the final destination.

Alternative Prefix field contains alternative prefixes assigned to the destination interface other than the one included in the Destination Address field of the IPv6 header [1].

The Alternative Prefix Extension Header is identified by a Next Header value of xx (TBD by IANA) in the immediately preceding header.

The position of the new extension header relative to the ones already defined is after the routing header and before the fragment header, since it is to be processed by intermediate routers when no route to destination is found.

The intended usage of the above Extension header is the following: when a router receives a packet and it has no route to the address contained in the destination field, the router must look for an Alternative Prefix Extension Header. If such header is included in the packet, and the value of Pleft is different than zero, then the router must swap the 64 most significant bits of the Destination address with the ones located in the position number (Ext Hdr Len minus Pleft) of the extension header. Then the router must update

the Pleft by Pleft minus one. Finally, the router must try to

forward the packet to the new destination address. If there is no Alternative Prefix Extension Header or the Pleft value is zero, then the packet must be discarded and an ICMP error must be sent to the source.

```
If (No Route to Destination) AND (Exists Alternative Prefix Extension Header)
then
 {
   if Pleft = 0 {
       Discard packet
                  }
   else {
      if Pleft is greater than Hdr Ext Len
         {
         send an ICMP Parameter Problem, Code 0, message to the Source
         Address, pointing to the Pleft field, and discard the
         packet
         }
      else {
         decrement Pleft by 1;
         compute i, the index of the next Prefix to be used by subtracting Pleft from Hdr Ext Len
         and swap the prefix of the IPv6 Destination Address and the Alternative Prefix #i
         resubmit the packet to the IPv6 module for transmission
         to the new destination
           }
       }
  }
```

3. Solution description

3.1 Scenario Description

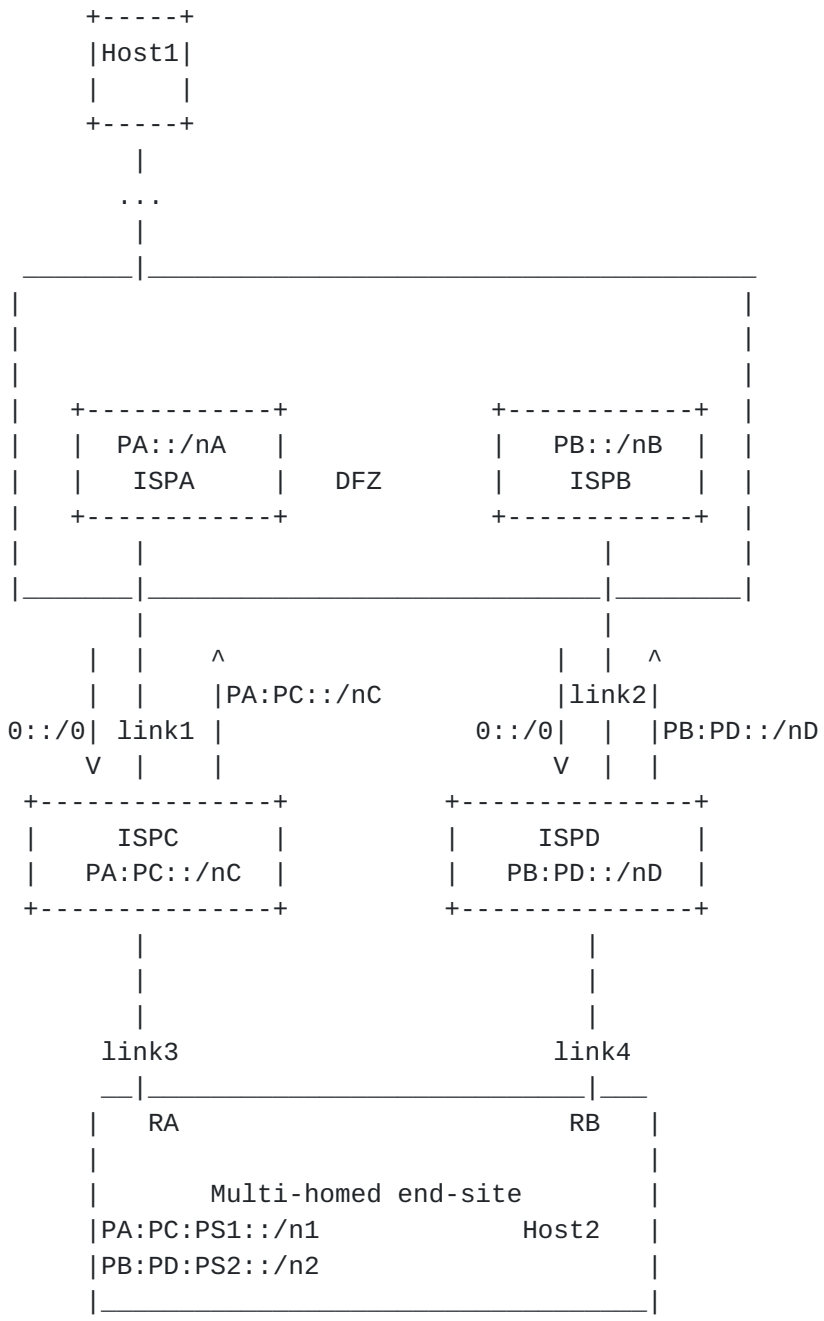

The considered scenario is as follows:

A Multi-homed end-site obtains global connectivity through two ISPs

i.e. ISPC and ISPD. These ISPs do not belong to the Default free

zone and they buy transit from ISPA and ISPB respectively. ISPA and ISPB do belong to the Default Free Zone i.e. at least one of their routers have full routing information. In the figure above, some of the routing information exchanged between peers is included. Since the end-site is multi-homed, it has obtained two address ranges: one delegated from ISPC address range i.e. PA:PC:PS1::/n1 and the other one from ISPD address space i.e. PB:PD:PS2::/n2. ISPC and ISPD have obtained a range of the address space from the address range assigned to their respective providers, i.e. ISPA and ISPB. So ISPA has delegated the range PA:PC::/nC to ISPC and ISPB has delegated the range PB:PD::/nD to ISPD.

3.2 Normal operation.

Let´s now consider the possible communication between Host1 (a given host in the Internet) and Host2 (a host belonging to the multi-homed end-site considered)

Since Host2 belongs to the Site, it has at least two addresses i.e. PA:PC:PS1:PL1:IIdHost2 and PB:PD:PS2:PL2:IIdHost2, which will be included in the DNS (if we suppose that Host2 wants to be reached through the two providers). It should be noted that the solution requires that all addresses of the same interface used in the solution share the Interface identifier part

Communication initiated by Host2.

Host2 sends a packet to Host1 address and it includes a Alternative Prefix Destination option with all the different prefixes it is willing to use to receive replies to this packet. When Host1 replies, it addresses the packet to the source address included in the first packet and it also includes in the reply packet an Alternative Prefix Extension Header with the prefixes included in the Alternative Prefix Destination option of the initial packet. When Host2 receives the reply, it verifies that the destination address and all the prefixes included in the Alternative Prefix Extension Header are assigned to its interfaces. If at least one of the derived addresses is not assigned to any of the interfaces, the packet is discarded (See Security Considerations Section). Even if different packets of a given communication may have different destination addresses, Host2 must present them to its upper layer as if they had the same destination address. This can be done since it is possible to identify the original destination address used by Host1 in the following way: If the Ext Hdr Len value in the Alternative Prefix Extension Header is equal to the value of the Pleft field, then the original Destination address is the one included in the Destination Address field of the IPv6 header. If the

Ext Hdr Len value in the Alternative Prefix Extension Header is

greater than the value of the Pleft field, then the original Destination address can be reconstructed by replacing the prefix of the address included in the destination address field of the IPv6 header by the first prefix included in the Extension header. Then the packet exchange will continue as above.

Communication initiated by Host1:

Host1 performs an A-type query to the DNS and it obtains two addresses i.e. PA:PC:PS1:PL1:IIdHost2 and PB:PD:PS2:PL2:IIdHost2.

At this point Host1 can make two different uses of the obtained information:

First option: Host1 uses one of the obtained addresses as the destination address and it includes the other address in an Alternative Prefix Extension Header. This option would provide the same treatment for all the packets sent by Host1 and in particular it would provide fault tolerance for this packet. However, this option would imply some changes in the way applications manage multiple addresses obtained from a DNS query.

Second option: The first packet is sent using available fault tolerance capabilities when multiple addresses are available i.e. Host1 sends a first packet with one of the obtained addresses and if no reply is obtained it retries with an alternative address. When finally a reply is received, an Alternative Prefix Destination Option is included in it, so that alternative addresses are learned, as in the previous case.

Eventually, in either case, packets flowing from Host1 to Host2 will carry the Alternative Prefix Extension Header, and communication will continue as detailed above.

3.3 Fault Tolerance Support

We will next study fault tolerance performance of the solution. Let´s suppose that Host1 is sending packets to Host2 address PA:PC:PS1:PL1:IIdHost2 and Link1 fails. In this case, when next packets arrive to ISPA, there will be no route to the destination, so the ISPA router with no route to destination in its routing tables will look for an Alternative Prefix Extension Header in the packet. If this header is found, it will be processed and the prefix of the destination address will be replaced with the one found in the extension header, and the packet will follow the alternative route towards its destination. Some may argue that Alternative Prefix Extension Header processing imposes an unacceptable load in routers, specially in Core Routers. Another issue that could be raised, is

the need for upgrading all the routers of the ISP in order to be able to process the newly defined Extension Header. A workaround for this issues can be found by noting that the extension header processing can be performed by specific upgraded routers connected to the ISP network which would work in the following way: These upgraded routers announce a default route (in our example, the upgraded router is connected to the ISPA network and it announces the a route to 0/0). Then if link1 is working properly, longest prefix match rule will make packets flow through link1. If link1 is down, packets will be forwarded to the upgraded router, that will process the Alternative Prefix Extension Header, swapping prefix information. Once this is done, it will forward the packet to the ISPA network, and then to the alternative route. A slightly different approach is needed to provide a sink route for packets with unreachable destination address when link3 fails. Since ISPC obtains a default route from its provider ISPA, it is not possible to announce a default route to sink packets with unreachable destination, as in the previous case where the ISP (ISPA) belongs to the default free zone. In this case, the upgraded routers announce a route to the address range allocated to the ISP (in our example, the upgraded router is connected to the ISPC network and it announces the a route to PA:PC::/nC). Then if link3 is working properly, the longest prefix match rule will make packets flow through link3. If link3 is down, packets will be forwarded to the upgraded router, who will process the Alternative Prefix Extension Header, swapping prefix information. Once this is done, it will forward the packet to the ISPC network, and then to the alternative route.

4. Security Considerations.

The extension header and the destination option defined above may seem to introduce new security risks, since they seem to enable the inclusion of spoofed alternative address. This would allow different type of attacks such as communication hijacking. However, this situation can detected by the host belonging to the multi-homed site, since if any of the addresses included in the Alternative Prefix Extension Header does not correspond to a configured one, the packet will be discarded. This makes us to conclude that packets carrying the newly defined option or header are not more susceptible to attacks than regular unicast packets. It must be noted that both types of packets are susceptible to man in the middle attacks, but the goal of this solution is not improving security features but avoiding the introduction of new security risks.

IPSec support: Alternative Prefix Destination option does not change in route so interaction with IPSec is straightforward. Alternative Prefix Extension Header can be modified en-route, as well as the destination address of the IPv6 header during extension header processing. However, original IPv6 header and extension header can be reconstructed at the destination with the information included in the packet, so this solution is compatible with IPSec.

5. Acknowledgements

We would like to thank Ignacio Soto, Juan Francisco Rodriguez Hervella, Iljitsch van Beijnum and Michael Py for their reviews and comments.

References

[1] Hinden, R. and S. Deering, "Internet Protocol, Version 6 (IPv6) Specification", December 1998.

[2] Hinden, R. and S. Deering, "IP version 6 Addressing Architecture", July 1998.

Authors' Addresses

Marcelo Bagnulo
Universidad Carlos III de Madrid
Av. Universidad 30
Leganes, Madrid 28911
SPAIN

Phone: 34 91 6249500
EMail: marcelo@it.uc3m.es
URI: http://www.it.uc3m.es/marcelo

Alberto Garcia-Martinez
Universidad Carlos III de Madrid
Av. Universidad 30
Leganes, Madrid 28911
SPAIN

Phone: 34 91 6249500
EMail: alberto@it.uc3m.es
URI: http://www.it.uc3m.es/alberto

Full Copyright Statement

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.