Network Working Group
Internet-Draft
Expires: August 26, 2003

M. Bagnulo
A. Garcia-Martinez
I. Soto
UC3M

February 25, 2003

Application of the MIPv6 protocol to the multi-homing problem
draft-bagnulo-multi6-mnm-00

Status of this Memo

This Internet-Draft will expire on August 26, 2003.

Copyright Notice

Abstract

This note attempts to describe how to apply the MIPv6 protocol to provide fault tolerance to transport layer connections established between a multi-homed host and other hosts in the Internet. Specifically, this note addresses the usage of MIPv6 signaling messages to convey information about alternative address to be used when an outage occurs. Additionally, possible mechanisms to detect failures affecting the currently used path are explored.

1. Introduction

Several times it has been claimed that the MIPv6 [1] protocol could be a useful tool to deal with the multi-homing problem. A few years ago, the application of MIPv6 protocol to the multi-homing problem was proposed by F. Dupont in [2]. Since that time, the MIPv6 protocol has been greatly improved and substantial modifications have been introduced, particularly in the security aspects of the protocol. Recently, C. Huitema suggested in [3] that mobility extensions could be used to convey alternative address information of multi-homed hosts. This note attempts to complement previous work by providing a complete proposal of how to apply the MIPv6 protocol to provide fault tolerance to transport layer connections established between a multi-homed host and hosts in the Internet. Specifically, this note addresses the usage of MIPv6 signaling messages to convey information about alternative address to be used when an outage occurs. Additionally, possible mechanisms to detect failures affecting the currently used path are explored.

2. Acronyms

MHH: Multi-Homed Host

CN: Correspondent Node

BU: Binding Update

BA: Binding Acknowledgment

HoA: Home Address

CoA: Care-of Address

HoTI: Home Test Init

HoT: Home Test

CoTI: Care-of Test Init

CoT: Care-of Test

PA: Provider Aggregatable

SEAA: Site Exit Anycast Address

SEAAA: Site Exit Anycast Address corresponding to ISPA

SEAAB: Site Exit Anycast Address corresponding to ISPB

3. Application Scenario

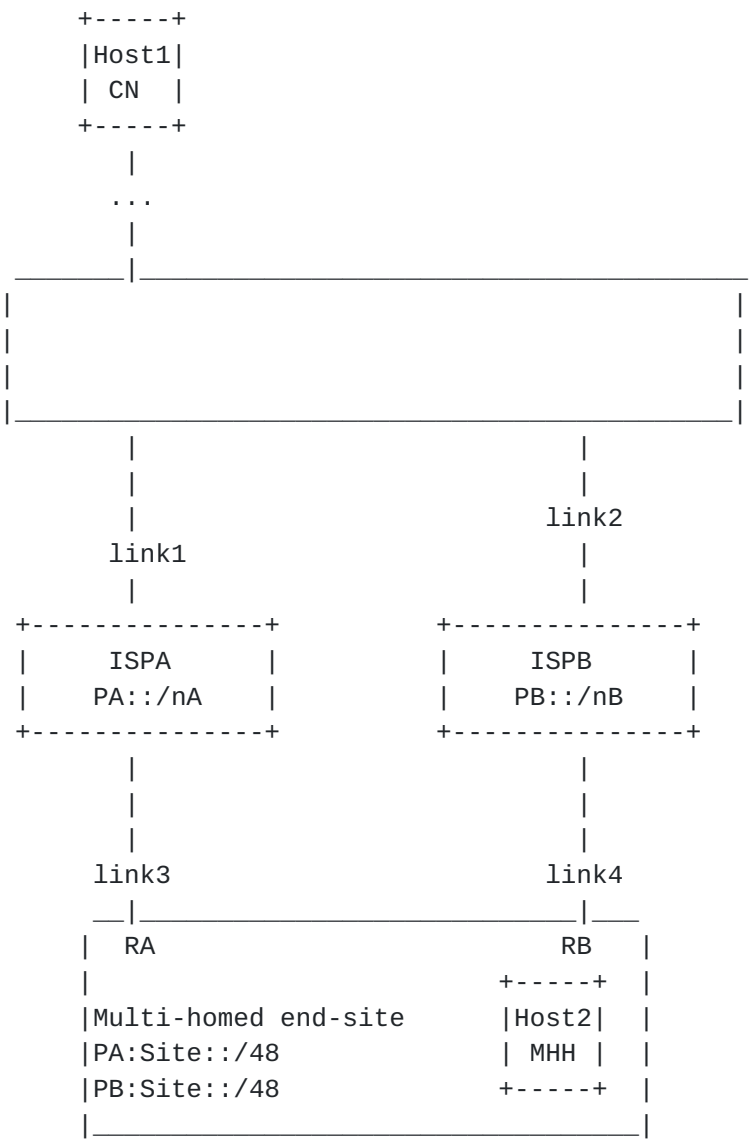

The application scenario consists of a multi-homed end-site that obtains global connectivity through two (or more) ISPs i.e. ISPA and ISPB. Since the end-site is multi-homed and provider aggregatable addresses are being used, the site has obtained two address ranges: one delegated from ISPA address range i.e. PA:Site::/48 and the other one from ISPB address space i.e. PB:Site::/48. Furthermore, in order to benefit from multi-homing, hosts within the site have to be

reachable through both ISPs. This implies that hosts have to configure one address from each ISP address range that the site has obtained. For instance, in the case of Host2, two addresses are configured i.e. PA:Site:Host2 and PB:Site:Host2. So, this configuration provides some Fault Tolerance capabilities since Host1 can reach Host2 through ISPA, using as destination address PA:Site:Host2 and it can also reach it through ISPB using as destination address PB:Site:Host2. This means that if there is an outage in one of the ISPs, ISPA for instance, Host1 can still reach Host2 using the alternative address i.e. PB:Site:Host2. However, this configuration does not allow the preservation of established connections through an outage event. This is the result of following constraints:

- Most connections established at the transport level and above identify the endpoints involved in the communication by their IP addresses, imposing that they must remain unchanged during the lifetime of the connection.

- In order to preserve aggregation benefits, Provider Aggregatable addresses delegated to the end-site by an ISP are to be routed toward the site through this ISP, so that it routes them toward the final destination. In the application scenario considered, addresses containing the prefix PA::/nA are routed to ISPA, who then routes them to the end-site considered.

These constraints imply that if a connection between Host1 and Host2 is established using PA:Site:Host2 as address for Host2, packets flowing to Host2 will be routed through ISPA and only through ISPA. Then, if during the lifetime of this connection an outage occurs in ISPA, the connection will be dropped, even if a path between Host1 and Host2 is available. This is so because packets whose final destination address contains the PA::/nA prefix have no available route to Host2, since they can not be routed through ISPB and packets addressed to the alternative destination of Host2 (PB:Site:Host2) are not recognized by transport and uppers layers as belonging to the connection established using PA:Site:Host2. This note presents a mechanism for preserving established communications during an outage based in the MIPv6 protocol.

4. Application of the MIPv6 protocol to the multi-homing problem

4.1 Required capabilities

In order to preserve established connections throughout an outage, the following capabilities are required:

1- A path failure detection mechanism, that enables end-hosts to detect outages in the path that is currently being used. When a failure is detected a recovery mechanism, such as routing packets through an alternative path, is triggered.

2- A protocol to inform the other end of the communication about the alternative path that is to be used. Since Provider Aggregatable address are used, alternative paths (alternative ISPs) are represented by alternative destination addresses. So the protocol is used for conveying alternative destination address.

3- A mechanism that allows packets carrying the alternative address as destination address to be recognized as belonging to the established connection. In order to be transparent to transport layer and above, such mechanism must restore the original destination address when the final destination is reached.

4- Tools to ensure compatibility with ingress filtering mechanisms. Since an alternative ISP will be used when a outage occurs, packets carrying the original source address would be incompatible with ingress filtering mechanisms.

4.2 Overview

MIPv6 protocol provides the tools needed to satisfy all the above requirements, as it will presented next.

In order to apply the MIPv6 protocol to the considered scenario, the first step is to map the multi-homing scenario to a mobility scenario. Since the multi-homed host (MHH) has the need to use multiple alternative addresses in a given connection, it will have the role of mobile node, and the node that it is communicating with will have the role of Correspondent Node (CN). It is assumed that Correspondent Nodes have support for route optimization. Home Agent capabilities are not required.

4.2.1 Required capability #2

The most natural application of the MIPv6 protocol seems to be the usage of Binding Update (BU) messages to inform the Correspondent node that an alternative address is to be used for the established

communication, fulfilling requirement #2. So, when the MHH detects that the currently used path becomes unavailable, it would send a Binding Update message to the Correspondent Node, informing that an alternative address is to be used. In MIPv6 terminology, the original address would be the Home Address (HoA) and the new alternative address would be the Care-of Address (CoA). However, MIPv6 security requirements impose the return routability procedure to enable the required BU message authorization. Such procedure implies the exchange of Home Test Init (HoTI) and Home Test (HoT) messages using the HoA and the exchange of Care-of Test Init (CoTI) and Care-of Test (CoT) messages using the CoA. Such exchanges are designed to verify that the host reachable through both the CoA and the HoA is the same. This means that the MHH needs to be reachable through both paths when these exchanges are performed, implying that these exchanges can not be performed successfully once an outage has occurred. So, the return routability procedure should be performed when a connection with a new CN is established allowing that this connection is protected during its complete lifetime.

However, nonces used for the generation of the home keygen token and the care-of keygen token have a limited lifetime, imposing periodical return routability checks, in order to ensure that valid BU authorization information is available when an outage occurs. The time constraints imposed by MIPv6 are:

1- It is recommended by MIPv6 specification that nonces remain valid for at least MAX_TOKEN_LIFE seconds i.e. 210 seconds after it has been used to construct a return routability message.

2- MIPv6 also specifies that the CN must not accept nonces beyond MAX_NONCE_LIFE seconds i.e. 240 seconds after their first use.

These constraints impose the performance of the the return routability procedure every MAX_TOKEN_LIFE minus the time required to perform the procedure, which would include the Round Trip Time (RTT) and the processing time. If the typical value used for TCP connection establishment timeout (75 seconds) is accepted as a reasonable upper bound to the RTT, and the processing time is considered to be negligible compared to the RTT considered, the return routability procedure needs to be performed every 135 seconds.

4.2.2 Required capability #3

Once that the availability of the information needed to authorize BU messages is guaranteed, the MHH is prepared to re-route its connections through an alternative address when an outage occurs. So, when the outage is detected, the MHH will send a BU message to the CN, informing that a new address (CoA) is to be used. Then, the CN

will address packets to the new CoA. However, packets addressed to CoA have to be recognized as belonging to the established connection. This can be achieved by using the Type 2 Routing Header specified in MIPv6, in the same way that it is used for supporting mobility. That is, after receiving a valid BU, the CN addresses packets to the MHH to the new CoA, and it also includes a Type 2 Routing Header carrying the HoA. The resulting behavior is that when these packets reach the MHH through the CoA, the Routing Header is processed and the HoA is restored and packets are presented to the transport and above layers as being addressed to the HoA, preserving established connections.

4.2.3 Required capability #1

Additionally, a failure detection mechanism that triggers the generation of BU messages is required to provide a complete solution. A possible approach for this is to rely on transport and/or application layer capabilities. Some transport protocols such as TCP provide a reliable service, implementing time-out and retransmission of packets. When unreliable transport protocols are used, some applications provide recovery mechanisms that imply retransmission of lost packets. These retransmission events can be used as a failure indication to trigger the usage of an alternative address. However, this approach requires that the transport layer and/or applications inform the IP layer about retransmission events, imposing modifications to current implementations. Besides, some applications, such as interactive voice applications, do not employ packet recovery mechanisms. In these cases, an additional failure detection mechanism has to be provided, so that these applications can benefit from multi-homing.

It is then deemed necessary to provide a failure detection mechanism. Such mechanism should be provided at the IP layer so that it is available for all applications and transport layers.

An simple end-host path failure detection mechanism can be based on the exchange of keep-alive messages. Since this is a network layer mechanism, a possibility is the exchange of ICMP messages. For instance, ICMP echo request/reply can be used, profiting that most IP stacks already include ICMP functionality. This would imply that only MHH stacks need to be modified in order to provide the failure detection mechanism.

Another possibility is the usage of HoTI/HoT messages. As it was mentioned above, return routability procedure needs to be performed periodically, implying message exchange between the MHH and the CN. However, in order to provide a failure detection mechanism, the message exchange frequency has to be increased not only because its period of 135 seconds may be deemed as unacceptable for certain

applications, but because valid authorization information is required for sending the BU message. Since a failure is indicated by at least one keep-alive message lost, it is necessary that after such event valid BU authorization information is still available, which implies that the information acquired during the previous message exchange is still valid. Then, assuming that a failure is indicated by the lost of two consecutive keep-alive packets, HoTI messages have to be generated by the MHH every MAX_TOKEN_LIFE/3 seconds i.e. 70 seconds. Then if two HoTI messages are lost, that is, no reply is received 140 seconds after a HoTI was sent, a BU message is generated and an alternative route is used. Note that HoT replies are linked to HoTI requests by the home init cookie parameter. The simple mechanism presented provides the minimum required functionality while respecting the timing imposed by the return routability procedure parameters. Failure response time can be improved by increasing the message exchange frequency. Moreover, adaptive mechanism, such as the TCP time-out calculation mechanism can also be contemplated, as long as they respect the timing constraints imposed by MIPv6. However, it should be noted that these changes only need to be performed at the MHH, since the CN role is limited to reply messages generated by the MHH. This allows that multiple failure detection policies can be implemented without affecting interoperability. It should also be noted that only HoTI/HoT message exchange frequency need to be increased, since only the currently used path need to be tested. In the case where more than two paths are available, testing all the available paths may provide some valuable information at the time an outage occurs, since it would enable a more educated path selection. In this case, CoTI/CoT messages can be used to test alternative paths and compare them. However, in the case of dual homed sites, where only one alternative paths is available, this testing does not seem to provide relevant information.

4.2.4 Required capability #4

When a MHH has multiple PA addresses configured in its interface, source address selection implies the selection of the ISP to be used in the return path. Moreover, because of ISP ingress filtering mechanism, source address selection also imposes the ISP to be used in the forward path, requiring additional functionalities at the multi-homed site to guarantee the appropriate ISP selection as discussed in [3]. Besides, when host based path failure detection mechanisms are used, the only party that has the information needed for selecting the path to be used is the host itself. So, in order to guarantee the compatibility with ingress filtering mechanisms, the MHH can select the exit ISP by means of a Routing Header. In order to simplify ISP selection, the Site Exit Anycast Address defined in [3] can be used. Then, after performing source address selection, the MHH addresses packets to the Site Exit Router Anycast Address

corresponding to the ISP that has assigned the address used as source address and it includes the final destination address in a Routing Header. It should be noted that the Site Exit Anycast Address is automatically deduced from the source address, so no additional configuration is required.

Additional complexity results when an outage occurs. In this case, an alternative ISP is to be used for coursing packets. Source address filtering mechanisms of the alternative ISP precludes the flow of packets carrying the address originally used i.e. the HoA. However, the CN only recognizes packets as belonging to the established connection if they carry the original HoA. In order to overcome this issue, the Home Address Destination Option is to be used, so that the source address corresponding to the alternative ISP (i.e. the CoA) is carried in the Source Address field of the IPv6 header and the original address (i.e. the HoA) is carried within the Home Address Option. When the packet is received by the CN, it processes the Home Address Option and restores the HoA as the Source Address. Note that when including the Routing Header to perform exit ISP selection, Site Exit Anycast Address have to be selected according to the source address actually carried in the Source Address field of the IPv6 packet and has no relation with the HoA carried in a Home Address Option.

4.3 Resulting Behavior

In this section, the complete operation of the solution in the application scenario is described.

First of all, a communication is established between the MHH and the CN. This communication can be initiated by any of the parties.

If the communication is initiated by the CN, it will first obtain at least one of the MHH's addresses, for instance using the DNS. If all the MHH's addresses are listed in the DNS, the CN will pick one and try to initiate the communication using this address. If a failure has occurred along the path, the attempt to initiate the communication will fail and the CN will try again with another address. Eventually, a packet from CN will reach MHH.

In the application scenario, Host1 obtains PA:Site:Host1 and PB:Site:Host1 from the DNS. Then it will try to initiate the communication using PA:Site:Host1 and if this fails it will try using PB:Site:Host1 (this is an arbitrary choice).

If the communication is initiated by the MHH, it will obtain CN's address, using for instance the DNS. Then, it will apply the source address selection algorithm which outcome will be the address to be

used as source address when sending packets to the CN. The MHH attempts to communicate with the CN using the selected source address. In order to avoid that the ISP ingress filtering mechanism discarded this packet, MHH addresses the packet to the Site Exit Anycast Address of the ISP that assigned the source address selected and includes the CN address within a Routing Header.If a failure along the path has occurred, the communication will fail and MHH will try with another source address, so that the packet is coursed through an alternative ISP.

In the application scenario, according to [3],

Site Exit Anycast Addresses for ISPA is
PA:Site:FFFF:FFFF:FFFF:FFFF:FFFF (Hereafter SEAAA)

Site Exit Anycast Addresses for ISPB is
PB:Site:FFFF:FFFF:FFFF:FFFF:FFFF (Hereafter SEAAB)

Then, MHH will try first to send a first packet with:

```
IPv6 Header
    Destination address: SEAAA
    Source address: PA:Site:Host2
Routing Header
    Type: 1
    Address#1: Host1
```

If this packet fails, it will try with:

```
IPv6 Header
    Destination address: SEAAB
    Source address: PB:Site:Host2
Routing Header
    Type: 1
    Address#1: Host1
```

When the packet arrives to the corresponding site exit router, the Routing Header will be processed and the Destination Address will be set to Host1.

Once that MHH starts sending packets to CN, different address roles

have been set i.e. the address used as Source Address in the first packet flowing from MHH to CN will be the HoA and the other available addresses of MHH will be CoAs. It should be noted that these roles are assigned when the communication is established and they are not preassigned. This means that any address can be HoA or CoA. Moreover, a given address could be used as HoA in a communication with a given host and used as CoA in a communication with another one.

In the application scenario, we suppose that the first packet flowing from MHH to CN has PA:Site:Host2 as source address, so that:

PA:Site:Host2 is HoA and PB:Site:Host2 is CoA

Once that the first packet is carried from MHH to CN, the MHH has to perform the return routability procedure in order to obtain valid authorization data. This data is to be used if an outage occurs to course packets using an alternative address. The return routability procedure consists in exchanging HoTI/HoT messages using the HoA and CoTI/CoT messages using CoA. These messages also have to include a Routing Header to select appropriate exit ISP.

In the application scenario the following messages are exchanged:

First MHH sends HoTI and CoTI messages:

```
HoTI message
IPv6 Header
    Destination Address: SEAAA
    Source Address: PA:Site:Host2
Routing Header
    Type: 1
    Address#1: Host1
Mobility Header
    Type: HoTI
    Home Init Cookie: HCookie
```

```
CoTI message
IPv6 Header
    Destination Address: SEAAB
    Source Address: PB:Site:Host2
Routing Header
    Type: 1
    Address #1: Host1
Mobility Header
    Type: CoTI
    Care-of Init Cookie: CCookie
```

The CN replies sending HoT and CoT messages as follows:

```
HoT message
IPv6 Header
    Destination Address: PA:Site:Host2
    Source Address: Host1
Mobility Header
    Type: HoT
    Home Init Cookie: HCookie (from HoTI message)
    Home Nonce Index: HNI
    Home Keygen Token: HKT
```

```
CoT message
IPv6 Header
    Destination Address: PB:Site:Host2
    Source Address: Host1
Mobility Header
    Type: CoT
    Care-of Init Cookie: CCookie (from CoTI message)
    Care-of Nonce Index: CNI
    Care-of Keygen Token: CKT
```

The goal of the CoTI/CoT message exchange is to maintain valid authorization data in case an outage occurs, so it is performed every 135 seconds.

HoTI/HoT message exchange has two goals: first it provides valid

authorization information in case an outage occurs and second it is used as a path failure detection mechanism. This second goal imposes that the HoTI/HoT exchange has to be performed every 70 seconds. Note that HoTI-HoT message correlation is detected using the Home Init cookie value.

If no outage occurs, the communication continues as it is, and HoTI/HoT and CoTI/CoT messages exchanges continue until the communication is finished.

The only difference with packets flowing from a single-homed site is that they carry a Routing Header to perform exit ISP selection.

Then, all packets flowing from the MHH to the CN carry the appropriate Routing Header to perform exit ISP selection according to the Source Address. In this case, packets carry PA:Site:Host2 as Source address so they carry SEAAA as Destination Address and they include Host1 in the Routing Header.

If an outage occurs, it will be detected by the failure detection mechanism and an alternative path will be used. If two consecutive packets HoTI are not replied within 140 seconds after the first message was sent, a failure will be assumed. This sets a timeout of 140 seconds for the first CoTI packet and a timeout of 70 secs for the second message. IF this is the case, a BU message is sent, informing the CN that the CoA will be used to exchange packets. This BU message will carry authorization information obtained through the last successful HoTI/HoT and CoTI/CoT message exchanges. Timing is such, that this information is still valid when the BU reaches CN. In order to confirm that the BU was successfully processed, a Binding Acknowledgment (BA) is requested. According to MIPv6 specification, if no BA is received within INITIAL_BINDACK_TIMEOUT (i.e. 1 second), MHH retransmits the BU message increasing the sequence number until a response is received. The retransmission uses an exponential back-off process, doubling the timeout every time the BU is retransmitted until a BA is received or the timeout period reaches MAX_BINDACK_TIMEOUT (i.e. 32 seconds). Retransmission of BU using this timeout can continue indefinitely according to the specification. However, in this particular case, the authorization information has a limited lifetime, so it is useless to continue with the retransmissions once the information is no longer valid, which occurs MAX_NONCE_LIFE (240 seconds) after the nonce was used for creating the home keygen token.

In the application scenario, suppose that two consecutive HoTI messages are not replied. In this case MHH send a BU message containing:

```
BU message
IPv6 Header
    Destination Address: SEAAB
    Source Address: PB:Site:Host2
Routing Header
    Type: 1
    Address #1: Host1
Destination Option Extension Header
    Home Address Option
        Home Address: PA:Site:Host2
Mobility Header
    Type: Binding Update
    Acknowledge: set
    Home Registration: reset
    Link-Local Compatibility:
    Key Management Mobility Capability: ignored by CN
    Sequence #: S
    Lifetime: 0xffff
    Mobility Options
        Binding Authorization Data Option
        Authenticator: First (96, HMAC_SHA1(Kbm, Mobility Data))
          Being:
          Mobility Data = PB:Site:Host2 | Host1 | Mobility Header Data
          Kbm = SHA1 (HKT | CKT) (from HoT and CoT messages respectively)
    Nonces Index Option
        Home Nonce Index: HNI (from HoT message)
        Care-of Nonce Index: CNI (from CoT message)
```

CN replies sending a BA as follows

```
BA message
IPv6 Header
    Destination Address: PB:Site:Host2
    Source Address: Host1
Routing Header
    Type: 2
    Home Address: PA:Site:Host2
Mobility Header
    Type: Binding Acknowledgment
    Status: 0 (Binding Update Accepted)
    Key Management Mobility Capability: 0
    Lifetime: granted by CN
    Sequence #: S (from BU)
    Mobility Options
        Binding Authorization Data Option
        Authenticator: First (96, HMAC_SHA1(Kbm, Mobility Data))
          Being:
          Mobility Data = PB:Site:Host2 | Host1 | Mobility Header Data
```

Once that the BU and BA messages have been exchanged, alternative ISP will be used to course packets between the MHH and the CN.

Packets from the CN to the MHH will contain a Type 2 Routing Header in order to be routed through the alternative ISP:

Packets from CN to MHH

```
IPv6 Header
    Destination Address: PB:Site:Host2
    Source Address: Host1
Routing Header
    Type: 2
    Home Address: PA:Site:Host2
```

Packets from the MHH to the CN will carry the Home Address Destination Option and a Routing Header to select exit ISP:

Packets from MHH to CN

```
IPv6 Header
    Destination Address: SEAAB
    Source Address: PB:Site:Host2
Routing Header
    Type: 1
    Address #1: Host1
Destination Option Header
    Home Address Option: PA:Site:Host2
```

The communication can continue using this route while the binding established at the CN remains valid. MIPv6 specification states that a binding established with CN using keys created using the return routability procedure must not exceed MAX_RR_BINDING_LIFE (i.e. 420 seconds). This is a major limitation for this application, since the binding can not be refreshed until the original route is restored and outages can last longer than 420 seconds. The possibility of increasing this value should be explored in order to adapt the protocol for this application.

Anyway, since the usage of the alternative route imposes additional overhead because of the Home Address Option and the Type 2 Routing Header, the original route is to be used as soon as it is restored. So the original route has to be probed in order to detect when it is restored. This can be done sending HoTI messages, so that when a HoT message is received, the original path can be restored.

Besides, it should be noted that no other alternative route can be used because no valid authorization data is available, since it can only be obtained through the exchange of HoTI/HoT messages through the original route. Because of this, using the failure detection mechanism to probe the alternative route does not seems to be relevant. Moreover, CoTI/CoT message exchange can be suspended until the original route is restored.

If a HoT message is finally received, a BU message is sent in order to restore the original route.

In the application scenario, MHH sends HoTI messages as follows

```
HoTI message
IPv6 Header
    Destination Address: SEAAA
    Source Address: PA:Site:Host2
Routing Header
    Type: 1
    Address#1: Host1
Mobility Header
    Type: HoTI
    Home Init Cookie: HCookie'
```

Until a HoT message is received from CN

```
HoT message
IPv6 Header
    Destination Address: PA:Site:Host2
    Source Address: Host1
Mobility Header
    Type: HoT
    Home Init Cookie: HCookie' (from HoTI message)
    Home Nonce Index: HNI'
    Home Keygen Token: HKT'
```

Then the MHH can send a BU message to restore the original route through ISPA.

```
BU message
IPv6 Header
    Destination Address: SEAAA
    Source Address: PA:Site:Host2
Routing Header
    Type: 1
    Address #1: Host1
Mobility Header
    Type: Binding Update
    Acknowledge: reset
    Home Registration: reset
    Link-Local Compatibility:
    Key Management Mobility Capability:
    Sequence #: S'
    Lifetime: 0 (indicates deleting a binding)
    Mobility Options
        Binding Authorization Data Option
        Authenticator: First (96, HMAC_SHA1(Kbm, Mobility Data))
          Being:
          Mobility Data = PA:Site:Host2 | Host1 | Mobility Header Data
          Kbm = SHA1 (HKT') (from HoT message)
    Nonces Index Option
        Home Nonce Index: HNI' (from HoT message)
        Care-of Nonce Index: 0
```

After receiving this message, the communication through the original route will be restored.

5. Evaluation of the solution

5.1 Limitations

The major limitation detected is that the lifetime of bindings established with the CN using keys created using the return routability procedure is limited to 420 seconds by the MIPv6 specification. This implies that the established connection is preserved for 7 minutes after the outage occurred. Since outages usually last longer than 7 minutes, this limitation drastically reduces the benefits provided by the solution. The possibility of extending the binding lifetime is to be explored if this solution is to be pursued. However, extending the binding lifetime may pose security problems, so this possibility has to be studied in great detail.

Another detected limitation is that once that a path has failed and another one is being used, it is not possible to switch the communication to a third path. This is so because no valid authorization information is available, since the return routability procedure requires exchange of information using the HoA, which is in this case unavailable.

Finally, a general limitation of the solutions based on hosts having multiple address is that it is difficult to implement policing. This is because hosts perform address selection and doing so they also select providers.

The two last limitations are mostly relevant for large/medium sites. So, they reduce the scope of the solution to small sites.

5.2 Benefits

The main benefit of the solution is its compatibility with the existent technology. For instance, changes needed are limited to hosts within the multi-homed site. Hosts in the outside of the site do not need to change almost anything in their implementation in order to communicate with a multi-homed site and benefit from multi-homing capabilities. The only change, if accepted, is the extension of maximum binding lifetime.

Besides, this solution is compatible with PA scheme granting the scalability of the routing system.

Finally, this solution does not requires that multi-homed sites to run BGP in contrast with many alternative solutions. This reduces the complexity of the solution and preserves the AS number space.

5.3 Possible optimizations

Multiple additional optimizations can be done to enhance the solution. Some are described next.

MIPv6 specification includes the possibility of piggybacking binding related messages in data packets as a future extension of the protocol. This would reduce the overhead imposed by the solution.

Other possible optimization that can be performed is related to the failure detection mechanism. The proposed mechanism provides minimum facilities. Improved algorithms can be proposed so that faster detection is provided. For instance adaptive mechanisms such as the ones used by TCP can be adopted. This document describes the constraints imposed by the MIPv6 specification. Any mechanism that honors these constraints is acceptable. Moreover, multiple mechanisms can be implemented in different hosts without compromising seamless interaction.

6. Security Considerations

The presented solution is based on the usage of the MIPv6 protocol, benefiting from MIPv6 security features. It should be assured by MIPv6 security experts that all the underlying assumptions of the mobility scenario remain valid.

Additionally, the possibility of extending the lifetime of bindings established with the CN using keys created using the return routability procedure may introduce security hazards that need to be carefully considered.

References

[1] Johnson, D., Perkins, C. and J. Arkko, "Mobility Support in IPv6", Internet Draft, Work in progress, May 2002.

[2] Dupont, F., "Multihomed routing domain issues for IPv6 aggregatable scheme", Internet Draft, Work in progress(Expired), March 2000.

[3] Huitema, C. and R. Draves, "Host-Centric IPv6 Multihoming", Internet Draft, Work in progress, June 2002.

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

Ignacio Soto
Universidad Carlos III de Madrid
Av. Universidad 30
Leganes, Madrid 28911
SPAIN

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

Intellectual Property Statement

Full Copyright Statement

Acknowledgement

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