

Network Working Group
Request for Comments: 4980
Category: Informational

C. Ng
Panasonic Singapore Labs
T. Ernst
INRIA
E. Paik
KT
M. Bagnulo
UC3M
October 2007

Analysis of Multihoming in Network Mobility Support

Status of This Memo

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

Abstract

This document is an analysis of multihoming in the context of network mobility (NEMO) in IPv6. As there are many situations in which mobile networks may be multihomed, a taxonomy is proposed to classify the possible configurations. The possible deployment scenarios of multihomed mobile networks are described together with the associated issues when network mobility is supported by RFC 3963 (NEMO Basic Support). Recommendations are offered on how to address these issues.

Table of Contents

1.	Introduction	3
2.	Classification	4
2.1.	(1,1,1): Single MR, Single HA, Single MNP	6
2.2.	(1,1,n): Single MR, Single HA, Multiple MNPs	6
2.3.	(1,n,1): Single MR, Multiple HAs, Single MNP	7
2.4.	(1,n,n): Single MR, Multiple HAs, Multiple MNPs	8
2.5.	(n,1,1): Multiple MRs, Single HA, Single MNP	8
2.6.	(n,1,n): Multiple MRs, Single HA, Multiple MNPs	9
2.7.	(n,n,1): Multiple MRs, Multiple HAs, Single MNP	9
2.8.	(n,n,n): Multiple MRs, Multiple HAs, Multiple MNPs	10

3.	Deployment Scenarios and Prerequisites	11
3.1.	Deployment Scenarios	11
3.2.	Prerequisites	13
4.	Multihoming Issues	14
4.1.	Fault Tolerance	14
4.1.1.	Failure Detection	15
4.1.2.	Path Exploration	16
4.1.3.	Path Selection	17
4.1.4.	Re-Homing	19
4.2.	Ingress Filtering	19
4.3.	HA Synchronization	21
4.4.	MR Synchronization	22
4.5.	Prefix Delegation	23
4.6.	Multiple Bindings/Registrations	23
4.7.	Source Address Selection	23
4.8.	Loop Prevention in Nested Mobile Networks	24
4.9.	Prefix Ownership	24
4.10.	Preference Settings	25
5.	Recommendations to the Working Group	26
6.	Conclusion	28
7.	Security Considerations	28
8.	Acknowledgments	29
9.	References	29
9.1.	Normative References	29
9.2.	Informative References	29
Appendix A.	Alternative Classifications Approach	32
A.1.	Ownership-Oriented Approach	32
A.1.1.	ISP Model	32
A.1.2.	Subscriber/Provider Model	33
A.2.	Problem-Oriented Approach	34
Appendix B.	Nested Tunneling for Fault Tolerance	35
B.1.	Detecting Presence of Alternate Routes	35
B.2.	Re-Establishment of Bi-Directional Tunnels	36
B.2.1.	Using Alternate Egress Interface	36
B.2.2.	Using Alternate Mobile Router	36
B.3.	To Avoid Tunneling Loop	37
B.4.	Points of Considerations	37

1. Introduction

The design goals and objectives of Network Mobility (NEMO) support in IPv6 are identified in [1], while the terminology is described in [2] and [3]. NEMO Basic Support (RFC 3963) [4] is the solution proposed by the NEMO Working Group to provide continuous Internet connectivity to nodes located in an IPv6 mobile network, e.g., like in an in-vehicle embedded IP network. The NEMO Basic Support solution does so by setting up bi-directional tunnels between the mobile routers (MRs) connecting the mobile network (NEMO) to the Internet and their respective home agents (HAs), much like how this is done in Mobile IPv6 [5], the solution for host mobility support. NEMO Basic Support is transparent to nodes located behind the MR (i.e., the mobile network nodes, or MNs), and as such, does not require MNs to take any action in the mobility management.

However, mobile networks are typically connected by means of wireless and thus less reliable links; there could also be many nodes behind the MR. A loss of connectivity or a failure to connect to the Internet has thus a more significant impact than for a single mobile node. Scenarios illustrated in [6] demonstrate that providing a permanent access to mobile networks typically require the use of several interfaces and technologies. For example, this is particularly useful for Intelligent Transport Systems (ITS) applications since vehicles are moving across distant geographical locations. Access would be provided through different access technologies (e.g., Wimax, Wifi, 3G) and through different access operators.

As specified in Section 5 of the NEMO Basic Support Requirements [1] (R.12), the NEMO WG must ensure that NEMO Basic Support does not prevent mobile networks to be multihomed, i.e., when there is more than one point of attachment between the mobile network and the Internet (see definitions in [3]). This arises either:

- o when an MR has multiple egress interfaces, or
- o the mobile network has multiple MRs, or
- o the mobile network is associated with multiple HAs, or
- o multiple global prefixes are available in the mobile network.

Using NEMO Basic Support, this would translate into having multiple bi-directional tunnels between the MR(s) and the corresponding HA(s), and may result in multiple Mobile Network Prefixes (MNs) available

to the MNs. However, NEMO Basic Support does not specify any particular mechanism to manage multiple bi-directional tunnels. The objectives of this memo are thus multifold:

- o to determine all the potential multihomed configurations for a NEMO, and then to identify which of these may be useful in a real-life scenario;
- o to capture issues that may prevent some multihomed configurations to be supported under the operation of NEMO Basic Support. It does not necessarily mean that the ones not supported will not work with NEMO Basic Support, as it may be up to the implementors to make it work (hopefully this memo will be helpful to these implementors);
- o to decide which issues are worth solving and to determine which WG is the most appropriate to address these;
- o to identify potential solutions to the previously identified issues.

In order to reach these objectives, a taxonomy for classifying the possible multihomed configurations is described in Section 2. Deployment scenarios, their benefits, and requirements to meet these benefits are illustrated in Section 3. Following this, the related issues are studied in Section 4. The issues are then summarized in a matrix for each of the deployment scenario, and recommendations are made on which of the issues should be worked on and where in Section 5. This memo concludes with an evaluation of NEMO Basic Support for multihomed configurations. Alternative classifications are outlined in the Appendix.

The readers should note that this document considers multihoming only from the point of view of an IPv6 environment. In order to understand this memo, the reader is expected to be familiar with the above cited documents, i.e., with the NEMO terminology as defined in [2] (Section 3) and [3], Motivations and Scenarios for Multihoming [6], Goals and Requirements of Network Mobility Support [1], and the NEMO Basic Support specification [4]. Goals and benefits of multihoming as discussed in [6], are applicable to fixed nodes, mobile nodes, fixed networks, and mobile networks.

2. Classification

As there are several configurations in which mobile networks are multihomed, there is a need to classify them into a clearly defined taxonomy. This can be done in various ways. A Configuration-Oriented taxonomy is described in this section. Two other

taxonomies, namely, the Ownership-Oriented Approach and the Problem-Oriented Approach, are outlined in Appendix A.1 and Appendix A.2.

Multihomed configurations can be classified depending on how many MRs are present, how many egress interfaces, Care-of Address (CoA), and Home Addresses (HoA) the MRs have, how many prefixes (MNPs) are available to the mobile network nodes, etc. We use three key parameters to differentiate the multihomed configurations. Using these parameters, each configuration is referred by the 3-tuple (x,y,z), where 'x', 'y', 'z' are defined as follows:

- o 'x' indicates the number of MRs where:

x=1 implies that a mobile network has only a single MR, presumably multihomed.

x=n implies that a mobile network has more than one MR.

- o 'y' indicates the number of HAS associated with the entire mobile network, where:

y=1 implies that a single HA is assigned to the mobile network.

y=n implies that multiple HAS are assigned to the mobile network.

- o 'z' indicates the number of MNPs available within the NEMO, where:

z=1 implies that a single MNP is available in the NEMO.

z=N implies that multiple MNPs are available in the NEMO.

It can be seen that the above three parameters are fairly orthogonal with one another. Thus, different values of 'x', 'y', and 'z' result in different combinations of the 3-tuple (x,y,z).

As will be described in the sub-sections below, a total of 8 possible configurations can be identified. One thing the reader has to keep in mind is that in each of the following 8 cases, the MR may be multihomed if either (i) multiple prefixes are available (on the home link, or on the foreign link), or (ii) the MR is equipped with multiple interfaces. In such a case, the MR would have multiple (HoA,CoA) pairs. Issues pertaining to a multihomed MR are also addressed in [7]. In addition, the readers should also keep in mind that when "MNP(s) is/are available in the NEMO", the MNP(s) may either be explicitly announced by the MR via router advertisement, or made available through Dynamic Host Configuration Protocol (DHCP) [8].

2.1. (1,1,1): Single MR, Single HA, Single MNP

The (1,1,1) configuration has only one MR, it is associated with a single HA, and a single MNP is available in the NEMO. The MR and the AR are connected to the Internet via a single Access Router (AR). To fall into a multihomed configuration, at least one of the following conditions must hold:

- o The MR has multiple interfaces and thus it has multiple CoAs;
- o Multiple global prefixes are available on the foreign link, and thus it has multiple CoAs; or
- o Multiple global prefixes are available on the home link, and thus the MR has more than one path to reach the HA.

Note that the case where multiple prefixes are available on the foreign link does not have any bearing on the MNPs. MNPs are independent of prefixes available on the link where the MR is attached to, thus prefixes available on the foreign link are not announced on the NEMO link. For the case where multiple prefixes are available on the home link, these are only announced on the NEMO link if the MR is configured to do so. In the present (1,1,1) configuration, only one MNP is announced.

A bi-directional tunnel would then be established between each (HoA,CoA) pair.

Regarding MNNs, they are (usually) not multihomed since they would configure a single global address from the single MNP available on the link they are attached to.




Figure 1: (1,1,1): 1 MR, 1 HA, 1 MNP

2.2. (1,1,n): Single MR, Single HA, Multiple MNPs

The (1,1,n) configuration has only one MR, it is associated with a single HA, and two or more MNPs are available in the NEMO.

The MR may itself be multihomed, as detailed in Section 2.1. In such a case, a bi-directional tunnel would be established between each (HoA,CoA) pair.

Regarding MNPs, they are multihomed because several MNPs are available on the link they are attached to. The MNPs would then configure a global address from each MNP available on the link.




Figure 2: (1,1,n): 1 MR, 1 HA, multiple MNPs

2.3. (1,n,1): Single MR, Multiple HAs, Single MNP

The (1,n,1) configuration has only one MR and a single MNP is available in the NEMO. The MR, however, is associated with multiple HAs.

The NEMO is multihomed since it has multiple HAs, but in addition, the conditions detailed in Section 2.1 may also hold for the MR. A bi-directional tunnel would then be established between each (HoA,CoA) pair.

Regarding MNPs, they are (usually) not multihomed since they would configure a single global address from the single MNP available on the link they are attached to.




Figure 3: (1,n,1): 1 MR, multiple HAs, 1 MNP

2.4. (1,n,n): Single MR, Multiple HAs, Multiple MNPs

The (1,n,n) configuration has only one MR. However, the MR is associated with multiple HAs and more than one MNP is available in the NEMO.

The MR is multihomed since it has multiple HAs, but in addition, the conditions detailed in Section 2.1 may also hold. A bi-directional tunnel would then be established between each (HoA,CoA) pair.

Regarding MNPs, they are multihomed because several MNPs are available on the link they are attached to. The MNPs would then configure a global address with each MNP available on the link.




Figure 4: (1,n,n): 1 MR, multiple HAs, multiple MNPs

2.5. (n,1,1): Multiple MRs, Single HA, Single MNP

The (n,1,1) configuration has more than one MR advertising global routes. However, the MR(s) are associated with a single HA, and there is a single MNP available in the NEMO.

The NEMO is multihomed since it has multiple MRs, but in addition the conditions detailed in Section 2.1 may also hold for each MR. A bi-directional tunnel would then be established between each (HoA,CoA) pair.

Regarding MNPs, they are (usually) not multihomed since they would configure a single global address from the single MNP available on the link they are attached to.




Figure 5: (n,1,1): Multiple MRs, 1 HA, 1 MNP

2.6. (n,1,n): Multiple MRs, Single HA, Multiple MNPs

The (n,1,n) configuration has more than one MR; multiple global routes are advertised by the MRs and multiple MNPs are available within the NEMO.

The NEMO is multihomed since it has multiple MRs, but in addition, the conditions detailed in Section 2.1 may also hold for each MR. A bi-directional tunnel would then be established between each (HoA,CoA) pair.

Regarding MNNs, they are multihomed because several MNPs are available on the link they are attached to. The MNNs would then configure a global address with each MNP available on the link.




Figure 6: (n,1,n): Multiple MRs, 1 HA, multiple MNPs

2.7. (n,n,1): Multiple MRs, Multiple HAs, Single MNP

The (n,n,1) configuration has more than one MR advertising multiple global routes. The mobile network is simultaneously associated with multiple HAs and a single MNP is available in the NEMO.

The NEMO is multihomed since it has multiple MRs and HAs, but in addition, the conditions detailed in Section 2.1 may also hold for each MR. A bi-directional tunnel would then be established between each (HoA,CoA) pair.

Regarding MNPs, they are (usually) not multihomed since they would configure a single global address from the single MNP available on the link they are attached to.




Figure 7: (n,n,1): Multiple MRs, Multiple HAs, 1 MNP

2.8. (n,n,n): Multiple MRs, Multiple HAs, Multiple MNPs

The (n,n,n) configuration has multiple MRs advertising different global routes. The mobile network is simultaneously associated with more than one HA and multiple MNPs are available in the NEMO.

The NEMO is multihomed since it has multiple MRs and HAs, but in addition, the conditions detailed in Section 2.1 may also hold for each MR. A bi-directional tunnel would then be established between each (HoA,CoA) pair.

Regarding MNPs, they are multihomed because several MNPs are available on the link they are attached to. The MNPs would then configure a global address with each MNP available on the link.




Figure 8: (n,n,n): Multiple MRs, HAs, and MNPs

3. Deployment Scenarios and Prerequisites

The following generic goals and benefits of multihoming are discussed in [6]:

1. Permanent and Ubiquitous Access
2. Reliability
3. Load Sharing
4. Load Balancing/Flow Distribution
5. Preference Settings
6. Aggregate Bandwidth

These benefits are now illustrated from a NEMO perspective with a typical instance scenario for each case in the taxonomy. We then discuss the prerequisites to fulfill these.

3.1. Deployment Scenarios

x=1: Multihomed mobile networks with a single MR

o Example 1:

MR with dual/multiple access interfaces (e.g., 802.11 and GPRS capabilities). This is a (1,1,*) if a single HA is used for both. If two independent HAs are used, this is a (1,n,n) configuration.

Benefits: Ubiquitous Access, Reliability, Load Sharing, Preference Settings, Aggregate Bandwidth.

x=n: Multihomed mobile networks with multiple MRs

- o Example 1:

Train with one MR in each car, all served by the same HA, thus a (n,1,*) configuration. Alternatively, the train company might use different HAs, in different countries, thus a (n,n,n) configuration.

Benefits: Ubiquitous Access, Reliability, Load Sharing, Aggregate Bandwidth.

- o Example 2:

Wireless personal area network with a GPRS-enabled phone and a WiFi-enabled PDA. This is a (n,n,n) configuration if different HAs are also used.

Benefits: Ubiquitous Access, Reliability, Preference Settings, Aggregate Bandwidth.

y=1: Multihomed mobile networks with a single HA

- o Example:

Most single HA cases in above examples.

y=n: Multihomed mobile networks with multiple HAs

- o Example 1:

Most multiple HAs cases in above examples.

- o Example 2:

Transatlantic flight with a HA in each continent. This is a (1,n,1) configuration if there is only one MR.

Benefits: Ubiquitous Access, Reliability, Preference Settings (reduced delay, shortest path).

z=1: Multihomed mobile networks with a single MNP

- o Example:

Most single HA cases in above examples.

z=n: Multihomed mobile networks with multiple MNPs

- o Example 1:

Most multiple HAS cases in above examples.

- o Example 2:

Car with a prefix taken from home (personal traffic is transmitted using this prefix and is paid by the owner) and one that belongs to the car manufacturer (maintenance traffic is paid by the car manufacturer). This will typically be a (1,1,n) or a (1,n,n,) configuration.

Benefits: Preference Settings

3.2. Prerequisites

In this section, requirements are stated in order to comply with the expected benefits of multihoming as detailed in [6].

At least one bi-directional tunnel must be available at any point in time between the mobile network and the fixed network to meet all expectations. But for most goals to be effective, multiple tunnels must be maintained simultaneously:

- o Permanent and Ubiquitous Access:

At least one bi-directional tunnel must be available at any point in time.

- o Reliability:

Both the inbound and outbound traffic must be transmitted or diverted over another bi-directional tunnel once a bi-directional tunnel is broken or disrupted. It should be noted that the provision of fault tolerance capabilities does not necessarily require the existence of multiple bi-directional tunnels simultaneously.

- o Load Sharing and Load Balancing:

Multiple tunnels must be maintained simultaneously.

- o Preference Settings:

Implicitly, multiple tunnels must be maintained simultaneously if preferences are set for deciding which of the available bi-directional tunnels should be used. To allow user/application to set the preference, a mechanism should be provided to the user/application for the notification of the availability of multiple bi-directional tunnels, and perhaps also to set preferences. A similar mechanism should also be provided to network administrators to manage preferences.

- o Aggregate Bandwidth:

Multiple tunnels must be maintained simultaneously in order to increase the total aggregated bandwidth available to the mobile network.

4. Multihoming Issues

As discussed in the previous section, multiple bi-directional tunnels need to be maintained either sequentially (e.g., for fault tolerance) or simultaneously (e.g., for load sharing).

In some cases, it may be necessary to divert packets from a (perhaps failed) bi-directional tunnel to an alternative (perhaps newly established) bi-directional tunnel (i.e., for matters of fault recovery, preferences), or to split traffic between multiple tunnels (load sharing, load balancing).

So, depending on the configuration under consideration, the issues discussed below may need to be addressed sometimes dynamically. For each issue, potential ways to solve the problem are investigated.

4.1. Fault Tolerance

One of the goals of multihoming is the provision of fault tolerance capabilities. In order to provide such features, a set of tasks need to be performed, including: failure detection, alternative available path exploration, path selection, and re-homing of established communications. These are also discussed in [9] by the Shim6 WG. In the following sub-sections, we look at these issues in the specific context of NEMO, rather than the general Shim6 perspective in [9]. In addition, in some scenarios, it may also be required to provide the mechanisms for coordination between different HAs (see Section 4.3) and also the coordination between different MRs (see Section 4.4).

4.1.1. Failure Detection

It is expected for faults to occur more readily at the edge of the network (i.e., the mobile nodes) due to the use of wireless connections. Efficient fault detection mechanisms are necessary to recover in timely fashion.

Depending on the NEMO configuration considered, the failure protection domain greatly varies. In some configurations, the protection domain provided by NEMO multihoming is limited to the links between the MR(s) and the HA(s). In other configurations, the protection domain allows to recover from failures in other parts of the path, so an end-to-end failure detection mechanism is required.

The failure detection capabilities required for each configuration are detailed below:

- o For the (1,1,*) cases, multiple paths are available between a single MR and a single HA. All the traffic to and from the NEMO flows through the MR and HA. Failure detection mechanisms need only to be executed between these two devices. This is a NEMO-/MIPv6-specific issue.
- o For the (n,1,*) cases, there is a single HA, so all the traffic to and from the NEMO will flow through it. The failure detection mechanisms need to be able to detect failure in the path between the used MR and the only HA. Hence, the failure detection mechanism needs only to involve the HA and the MRs. This is a NEMO/MIPv6 specific issue.
- o For the (n,n,*) cases, there are multiple paths between the different HAs and the different MRs. Moreover, the HAs may be located in different networks, and have different Internet access links. This implies that changing the HA used may not only allow recovering from failures in the link between the HA and the MR, but also from other failure modes, affecting other parts of the path. In this case, an end-to-end failure detection mechanism would provide additional protection. However, a higher number of failures is likely to occur in the link between the HA and the MR, so it may be reasonable to provide optimized failure detection mechanisms for this failure mode. The (n,n,n) case is hybrid, since selecting a different prefix results in a change of path. For this case, the Shim6 protocols (such as those discussed in [9]) may be useful.

Most of the above cases involve the detection of tunnel failures between HA(s) and MR(s). This is no different from the case of failure detection between a mobile host and its HA(s). As such, a

solution for MIPv6 should apply to NEMO as well. For case (n,*,*), an MR synchronization solution (see Section 4.4) should be able to complement a MIPv6 failure detection solution to achieve the desired functionality for NEMO.

In order for fault recovery to work, the MRs and HAS must first possess a means to detect failures:

- o On the MR's side, the MR can rely on router advertisements from access routers, or other layer-2 trigger mechanisms to detect faults, e.g., [10] and [11].
- o On the HA's side, it is more difficult to detect tunnel failures. For an ISP deployment model, the HAS and MRs can use proprietary methods (such as constant transmission of heartbeat signals) to detect failures and check tunnel liveness. In the subscriber model (see Appendix A.2: S/P model), a lack of standardized "tunnel liveness" protocol means that it is harder to detect failures.

A possible method is for the MRs to send binding updates more regularly with shorter Lifetime values. Similarly, HAS can return binding acknowledgment messages with smaller Lifetime values, thus forcing the MRs to send binding updates more frequently. These binding updates can be used to emulate "tunnel heartbeats". This, however, may lead to more traffic and processing overhead, since binding updates sent to HAS must be protected (and possibly encrypted) with security associations.

4.1.2. Path Exploration

Once a failure in the currently used path is detected, alternative paths have to be explored in order to identify an available one. This process is closely related to failure detection in the sense that paths being explored need to be alternative paths to the one that has failed. There are, however, subtle but significant differences between path exploration and failure detection. Failure detection occurs on the currently used path while path exploration occurs on the alternative paths (not on the one currently being used for exchanging packets). Although both path exploration and failure detection are likely to rely on a reachability or keepalive test exchange, failure detection also relies on other information, such as upper layer information (e.g., positive or negative feedback from TCP), lower layer information (e.g., an interface is down), and network layer information (e.g., as an address being deprecated or ICMP error message).

Basically, the same cases as in the analysis of the failure detection (Section 4.1.1) issue are identified:

- o For the (1,1,*) cases, multiple paths are available between a single MR and a single HA. The existing paths between the HA and the MR have to be explored to identify an available one. The mechanism involves only the HA and the MR. This is a NEMO-/MIPv6-specific issue.
- o For the (n,1,*) cases, there is a single HA, so all the traffic to and from the NEMO will flow through it. The available alternative paths are the different ones between the different MRs and the HA. The path-exploration mechanism only involves the HA and the MRs. This is a NEMO/MIPv6 specific issue.
- o For the (n,n,*) cases, there are multiple paths between the different HAs and the different MRs. In this case, alternative paths may be routed completely independent from one another. An end-to-end path-exploration mechanism would be able to discover if any of the end-to-end paths is available. The (n,n,1) case, however, seems to be pretty NEMO specific, because of the absence of multiple prefixes. The (n,n,n) case is hybrid, since selecting a different prefix results in a change of path. For this case, the Shim6 protocols (such as those discussed in [9]) may be useful.

Most of the above cases involve the path exploration of tunnels between HA(s) and MR(s). This is no different from the case of path exploration between a mobile host and its HA(s). As such, a solution for MIPv6 should apply to NEMO as well. For case (n,*,*), an MR synchronization solution (see Section 4.4) should be able to complement an MIPv6 path-exploration solution to achieve the desired functionality for NEMO.

In order to perform path exploration, it is sometimes also necessary for the MR to detect the availability of network media. This may be achieved using layer 2 triggers [10], or other mechanism developed/recommended by the Detecting Network Attachment (DNA) Working Group [11]. This is related to Section 4.1.1, since the ability to detect media availability would often imply the ability to detect media unavailability.

4.1.3. Path Selection

A path-selection mechanism is required to select among the multiple available paths. Depending on the NEMO multihoming configuration involved, the differences between the paths may affect only the part between the HA and the MR, or they may affect the full end-to-end

path. In addition, depending on the configuration, path selection may be performed by the HA(s), the MR(s), or the hosts themselves through address selection, as will be described in detail next.

The multiple available paths may differ on the tunnel between the MR and the HA used to carry traffic to/from the NEMO. In this case, path selection is performed by the MR and the intra-NEMO routing system for traffic flowing from the NEMO, and path selection is performed by the HA and intra-Home Network routing system for traffic flowing to the NEMO.

Alternatively, the multiple paths available may differ in more than just the tunnel between the MR and the HA, since the usage of different prefixes may result in using different providers; hence, in completely different paths between the involved endpoints. In this case, besides the mechanisms presented in the previous case, additional mechanisms for the end-to-end path selection may be needed. This mechanism may be closely related to source address selection mechanisms within the hosts, since selecting a given address implies selecting a given prefix, which is associated with a given ISP serving one of the home networks.

A dynamic path-selection mechanism is thus needed so that this path could be selected by:

- o The HA: it should be able to select the path based on some information recorded in the binding cache.
- o The MR: it should be able to select the path based on router advertisements received on both its egress interfaces or on its ingress interfaces for the (n,*,*) case.
- o The MNN: it should be able to select the path based on "Default Router Selection" (see [Section 6.3.6 Default Router Selection] [12]) in the (n,*,*) case or based on "Source Address Selection" in the (*,*,n) cases (see Section 4.7 of the present memo).
- o The user or the application: e.g., in case where a user wants to select a particular access technology among the available technologies for reasons, e.g., of cost or data rate.
- o A combination of any of the above: a hybrid mechanism should be also available, e.g., one in which the HA, the MR, and/or the MNNs are coordinated to select the path.

When multiple bi-directional tunnels are available and possibly used simultaneously, the mode of operation may be either primary-secondary (one tunnel is precedent over the others and used as the default

tunnel, while the other serves as a backup) or peer-to-peer (no tunnel has precedence over one another, they are used with the same priority). This questions which of the bi-directional tunnels would be selected, and based on which of the parameters (e.g., type of flow that goes into/out of the mobile network).

The mechanisms for the selection among the different tunnels between the MR(s) and the HA(s) seem to be quite NEMO/MIPv6 specific.

For (1,*,*) cases, they are no different from the case of path selection between a mobile host and its HA(s). As such, a solution for MIPv6 should apply to NEMO as well. For the (n,*,*) cases, an MR synchronization solution (see Section 4.4) should be able to complement an MIPv6 path-selection solution to achieve the desired functionality for NEMO.

The mechanisms for selecting among different end-to-end paths based on address selection are similar to the ones used in other multihoming scenarios, as those considered by Shim6 (e.g., [13]).

4.1.4. Re-Homing

After an outage has been detected and an available alternative path has been identified, a re-homing event takes place, diverting the existing communications from one path to the other. Similar to the previous items involved in this process, the re-homing procedure heavily varies depending on the NEMO multihoming configuration.

- o For the (*,*,1) configurations, the re-homing procedure involves only the MR(s) and the HA(s). The re-homing procedure may involve the exchange of additional BU messages. These mechanisms are shared between NEMO Basic Support and MIPv6.
- o For the (*,*,n) cases, in addition to the previous mechanisms, end-to-end mechanisms may be required. Such mechanisms may involve some form of end-to-end signaling or may simply rely on using different addresses for the communication. The involved mechanisms may be similar to those required for re-homing Shim6 communications (e.g., [13]).

4.2. Ingress Filtering

Ingress filtering mechanisms [14][15] may drop the outgoing packets when multiple bi-directional tunnels end up at different HAs. This could particularly occur if different MNPs are handled by different HAs. If a packet with a source address configured from a specific

MNP is tunneled to a HA that does not handle that specific MNP, the packet may be discarded either by the HA or by a border router in the home network.

The ingress filtering compatibility issue is heavily dependent on the particular NEMO multihoming configuration:

- o For the (*,*,1) cases, there is not such an issue, since there is a single MNP.
- o For the (1,1,*) and (n,1,1) cases, there is not such a problem, since there is a single HA, accepting all the MNPs.
- o For the (n,1,n) case, though ingress filtering would not occur at the HA, it may occur at the MRs, when each MR is handling different MNPs.
- o (*,n,n) are the cases where the ingress filtering presents some difficulties. In the (1,n,n) case, the problem is simplified because all the traffic to and from the NEMO is routed through a single MR. Such configuration allows the MR to properly route packets respecting the constraints imposed by ingress filtering. In this case, the single MR may face ingress filtering problems that a multihomed mobile node may face, as documented in [7]. The more complex case is the (n,n,n) case. A simplified case occurs when all the prefixes are accepted by all the HAs, so that no problems occur with the ingress filtering. However, this cannot be always assumed, resulting in the problem described below.

As an example of how this could happen, consider the deployment scenario illustrated in Figure 9: the mobile network has two mobile routers MR1 and MR2, with home agents HA1 and HA2, respectively. Two bi-directional tunnels are established between the two pairs. Each MR advertises a different MNP (P1 and P2 respectively). MNP P1 is registered to HA1, and MNP P2 is registered to HA2. Thus, MNPs should be free to auto-configure their addresses on any of P1 or P2. Ingress filtering could thus happen in two cases:

- o If the two tunnels are available, MNP cannot forward packet with source address equals P1.MNP to MR2. This would cause ingress filtering at HA2 to occur (or even at MR2). This is contrary to normal Neighbor Discovery [12] practice that an IPv6 node is free to choose any router as its default router regardless of the prefix it chooses to use.

- o If the tunnel to HA1 is broken, packets that would normally be sent through the tunnel to HA1 should be diverted through the tunnel to HA2. If HA2 (or some border router in HA2's domain) performs ingress filtering, packets with source address configured from MNP P1 may be discarded.

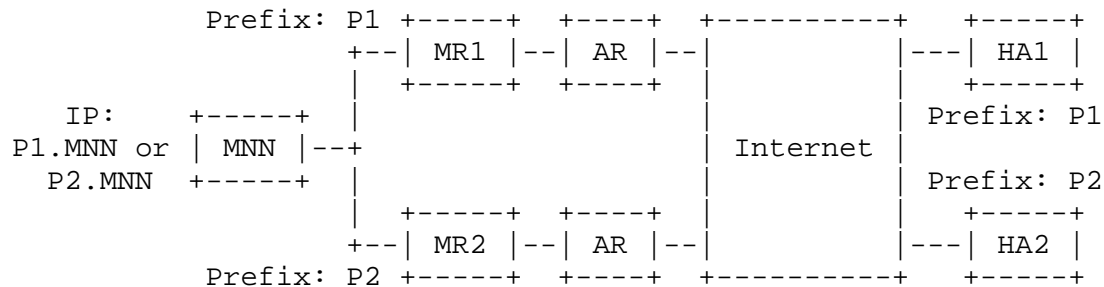


Figure 9: An (n,n,n) mobile network

Possible solutions to the ingress filtering incompatibility problem may be based on the following approaches:

- o Some form of source address-dependent routing, whether host-based and/or router-based where the prefix contained in the source address of the packet is considered when deciding which exit router to use when forwarding the packet.
- o The usage of nested tunnels for (*,n,n) cases. Appendix B describes one such approach.
- o Deprecating those prefixes associated to non-available exit routers.

The ingress filtering incompatibilities problems that appear in some NEMO multihoming configurations are similar to those considered in Shim6 (e.g., see [16]).

4.3. HA Synchronization

In the (*,n,*) configuration, a single MNP would be registered at different HAs. This gives rise to the following cases:

- o Only one HA may actively advertise a route to the MNP,
- o Multiple HAs at different domains may advertise a route to the same MNP.

This may pose a problem in the routing infrastructure as a whole if the HAs are located in different administrative domains. The implications of this aspect needs further exploration. A certain level of HA coordination may be required. A possible approach is to adopt an HA synchronization mechanism such as that described in [17] and [18]. Such synchronization might also be necessary in a (*,n,*) configuration, when an MR sends binding update messages to only one HA (instead of all HAs). In such cases, the binding update information might have to be synchronized between HAs. The mode of synchronization may be either primary-secondary or peer-to-peer. In addition, when a MNP is delegated to the MR (see Section 4.5), some level of coordination between the HAs may also be necessary.

This issue is a general mobility issue that will also have to be dealt with by Mobile IPv6 (see Section 6.2.3 in [7]) as well as NEMO Basic Support.

4.4. MR Synchronization

In the (n,*,*) configurations, there are common decisions that may require synchronization among different MRs [19], such as:

- o advertising the same MNP in the (n,*,1) configurations (see also "prefix delegation" in Section 4.5);
- o one MR relaying the advertisement of the MNP from another failed MR in the (n,*,n) configuration; and
- o relaying between MRs everything that needs to be relayed, such as data packets, creating a tunnel from the ingress interface, etc., in the (n,*,*) configuration.

However, there is no known standardized protocol for this kind of router-to-router synchronization. Without such synchronization, it may not be possible for a (n,*,*) configuration to achieve various multihoming goals, such as fault tolerance.

Such a synchronization mechanism can be primary-secondary (i.e., a master MR, with the other MRs as backup) or peer-to-peer (i.e., there is no clear administrative hierarchy between the MRs). The need for such mechanism is general in the sense that a multi-router site in the fixed network would require the same level of router synchronization.

Thus, this issue is not specific to NEMO Basic Support, though there is a more pressing need to develop an MR-to-MR synchronization scheme for proving fault tolerances and failure recovery in NEMO configurations due to the higher possibility of links failure.

In conclusion, it is recommended to investigate a generic solution to this issue although the solution would first have to be developed for NEMO deployments.

4.5. Prefix Delegation

In the (*,*,1) configurations, the same MNP must be advertised to the MNs through different paths. There is, however, no synchronization mechanism available to achieve this. Without a synchronization mechanism, MR may end up announcing incompatible MNPs. Particularly,

- o for the (*,n,1) cases, how can multiple HAS delegate the same MNP to the mobile network? For doing so, the HAS may be somehow configured to advertise the same MNP (see also "HA Synchronization" in Section 4.3).
- o for the (n,*,1) cases, how can multiple MRs be synchronized to advertise the same MNP down the NEMO-link? For doing so, the MRs may be somehow configured to advertise the same MNP (see also "MR Synchronization" in Section 4.4).

Prefix delegation mechanisms [20][21][22] could be used to ensure all routers advertise the same MNP. Their applicability to a multihomed mobile network should be considered.

4.6. Multiple Bindings/Registrations

When an MR is configured with multiple CoAs, it is often necessary for it to bind these CoAs to the same MNP.

This is a generic mobility issue, since Mobile IPv6 nodes face a similar problem. This issue is discussed in [7]. It is sufficient to note that solutions like [23] can solve this for both Mobile IPv6 and NEMO Basic Support. This issue is being dealt with in the Monami6 WG.

4.7. Source Address Selection

In the (*,*,n) configurations, MNs would be configured with multiple addresses. Source address selection mechanisms are needed to decide which address to choose from.

However, currently available source address selection mechanisms do not allow MNs to acquire sufficient information to select their source addresses intelligently (such as based on the traffic condition associated with the home network of each MNP). It may be desirable for MNs to be able to acquire "preference" information on each MNP from the MRs. This would allow default address selection

mechanisms, such as those specified in [24], to be used. Further exploration on setting such "preference" information in Router Advertisement based on performance of the bi-directional tunnel might prove to be useful. Note that source address selection may be closely related to path selection procedures (see Section 4.1.3) and re-homing techniques (see Section 4.1.4).

This is a general issue faced by any node when offered multiple prefixes.

4.8. Loop Prevention in Nested Mobile Networks

When a multihomed mobile network is nested within another mobile network, it can result in very complex topologies. For instance, a nested mobile network may be attached to two different root-MRs, thus the aggregated network no longer forms a simple tree structure. In such a situation, infinite loop within the mobile network may occur.

This problem is specific to NEMO Basic Support. However, at the time of writing, more research is recommended to assess the probability of loops occurring in a multihomed mobile network. For related work, see [25] for a mechanism to avoid loops in nested NEMO.

4.9. Prefix Ownership

When a $(n,*,1)$ network splits, (i.e., the two MRs split themselves up), MRs on distinct links may try to register the only available MNP. This cannot be allowed, as the HA has no way to know which node with an address configured from that MNP is attached to which MR. Some mechanism must be present for the MNP to either be forcibly removed from one (or all) MRs, or the implementors must not allow a $(n,*,1)$ network to split.

A possible approach to solving this problem is described in [26].

This problem is specific to NEMO Basic Support. However, it is unclear whether there is a sufficient deployment scenario for this problem to occur.

It is recommended that the NEMO WG standardize a solution to solve this problem if there is sufficient vendor/operator interest, or specify that the split of a $(n,*,1)$ network cannot be allowed without router renumbering.

4.10. Preference Settings

When a mobile network is multihomed, the MNNs may be able to benefit from this configuration, such as to choose among the available paths based on cost, transmission delays, bandwidth, etc. However, in some cases, such a choice is not made available to the MNNs.

Particularly:

- o In the (*,*,n) configuration, the MNNs can influence the path by source address selection (see Section 4.1.3 and Section 4.7).
- o In the (n,*,*) configuration, the MNNs can influence the path by default router selection (see Section 4.1.3).
- o In the (1,n,1) configuration, the MNNs cannot influence the path selection.

One aspect of preference setting is that the preference of the MNN (e.g., application or transport layer configuration) may not be the same as the preference used by MR. Thus, forwarding choices made by the MR may not be the best for a particular flow, and may even be detrimental to some transport control loops (i.e., the flow control algorithm for TCP may be messed up when MR unexpectedly performs load balancing on a TCP flow). A mechanism that allows the MNN to indicate its preference for a given traffic might be helpful here.

Another aspect of preference setting is that the MNN may not even be aware of the existence of multiple forwarding paths, e.g., the (1,n,1) configuration. A mechanism for the MR to advertise the availability of multiple tunneling paths would allow the MNN to take advantage of this, coupled with the previously mentioned mechanism that allows the MNN to indicate its preference for a given traffic.

This problem is general in the sense that IPv6 nodes may wish to influence the routing decision done by the upstream routers. Such a mechanism is currently being explored by various WGs, such as the NSIS and IPFIX WGs. It is also possible that the Shim6 layer in the MNNs may possess such a capability. It is recommended for vendors or operators to investigate into the solutions developed by these WGs when providing multihoming capabilities to mobile networks.

In addition, the Monami6 WG is currently developing a flow filtering solution for mobile nodes to indicate how flows should be forwarded by a filtering agent [27] (such as HA and mobile anchor points). It is recommended that the Monami6 WG consider the issues described here so that flow filtering can be performed by the MNN to indicate how flows should be forwarded by the MR.

5. Recommendations to the Working Group

Several issues that might impact the deployment of NEMO with multihoming capabilities were identified in Section 4. These are shown in the matrix below, for each of the eight multihoming configurations, together with indications from which WG(s) a solution to each issue is most likely to be found.

+=====+									
	# of MRs:	1	1	1	1	n	n	n	n
	# of HAs:	1	1	n	n	1	1	n	n
	# of Prefixes:	1	n	1	n	1	n	1	n
+=====+									
Fault Tolerance		*	*	*	*	*	*	*	*
+-----+									
Failure Detection		N/M	N/M	N/M	N/M	N/M	N/M	N	S
+-----+									
Path Exploration		N/M	N/M	N/M	N/M	N/M	N/M	N	S
+-----+									
Path Selection		N	S/M	M	S/M	N	S/N	N	S/N
+-----+									
Re-Homing		N/M	S	N/M	S	N/M	S	N/M	S
+-----+									
Ingress Filtering		.	.	.	t	.	.	.	N
+-----+									
HA Synchronization		.	.	N/M	N/M	.	.	N/M	N/M
+-----+									
MR Synchronization		G	G	G	G
+-----+									
Prefix Delegation		.	.	N	N	N	N	N	N
+-----+									
Multiple Binding/Registrations		M	M	M	M	M	M	M	M
+-----+									
Source Address Selection		.	G	.	G	.	G	.	G
+-----+									
Loop Prevention in Nested NEMO		N	N	N	N	N	N	N	N
+-----+									
Prefix Ownership		N	.	N	.
+-----+									
Preference Settings		G	G	G	G	G	G	G	G
+=====+									

N - NEMO Specific M - MIPv6 Specific G - Generic IPv6
 S - SHIM6 WG D - DNA WG
 . - Not an Issue t - trivial
 * - Fault Tolerance is a combination of Failure Detection, Path
 Exploration, Path Selection, and Re-Homing

Figure 10: Matrix of NEMO Multihoming Issues

The above matrix serves to identify which issues are NEMO-specific, and which are not. The readers are reminded that this matrix is a simplification of Section 4 as subtle details are not represented in Figure 10.

As can be seen from Figure 10, the following are some concerns that are specific to NEMO: Failure Detection, Path Exploration, Path Selection, Re-Homing, Ingress Filtering, HA Synchronization, Prefix Delegation, Loop Prevention in Nested NEMO, and Prefix Ownership. Based on the authors' best knowledge of the possible deployments of NEMO, it is recommended that:

- o A solution for Failure Detection, Path Exploration, Path Selection, and Re-Homing be solicited from other WGs.

Although Path Selection is reflected in Figure 10 as NEMO-Specific, the technical consideration of the problem is believed to be quite similar to the selection of multiple paths in mobile nodes. As such, we would recommend vendors to solicit a solution for these issues from other WGs in the IETF; for instance, the Monami6 or Shim6 WG.

- o Ingress Filtering on the (n,n,n) configuration can be solved by the NEMO WG.

This problem is clearly defined, and can be solved by the WG. Deployment of the (n,n,n) configuration can be envisioned on vehicles for mass transportation (such as buses, trains) where different service providers may install their own MRs on the vehicle/vessel.

It should be noted that the Shim6 WG may be developing a mechanism for overcoming ingress filtering in a more general sense. We thus recommend that the NEMO WG concentrate only on the (n,n,n) configuration should the WG decide to work on this issue.

- o A solution for HA Synchronization can be looked at in a mobility-specific WG, taking into consideration both mobile hosts operating Mobile IPv6 and MRs operating NEMO Basic Support.
- o A solution for Multiple Bindings/Registrations is presently being developed by the Monami6 WG.
- o Prefix Delegation should be reviewed and checked by the NEMO WG.

The proposed solutions [22] and [21] providing prefix delegation functionality to NEMO Basic Support should be reviewed in order to

make sure concerns, as discussed in Section 4.5, are properly handled.

- o Loop Prevention in Nested NEMO should be investigated.

Further research is recommended to assess the risk of having a loop in the nesting of multihomed mobile networks.

- o Prefix Ownership should be considered by the vendors and operators.

The problem of Prefix Ownership only occurs when a mobile network with multiple MRs and a single MNP can arbitrarily join and split. Vendors and operators of mobile networks are encouraged to input their views on the applicability of deploying such kind of mobile networks.

6. Conclusion

This memo presented an analysis of multihoming in the context of network mobility under the operation of NEMO Basic Support (RFC 3963). The purpose was to investigate issues related to such a bi-directional tunneling mechanism where mobile networks are multihomed and multiple bi-directional tunnels are established between Home Agent and Mobile Router pairs. For doing so, mobile networks were classified into a taxonomy comprising eight possible multihomed configurations. Issues were explained one by one and then summarized into a table showing the multihomed configurations where they apply, suggesting the most relevant IETF working group where they could be solved. This analysis will be helpful to extend the existing standards to support multihoming and to implementors of NEMO Basic Support and multihoming-related mechanisms.

7. Security Considerations

This is an informational document where the multihoming configurations under the operation of NEMO Basic Support are analyzed. Security considerations of these multihoming configurations, should they be different from those that concern NEMO Basic Support, must be considered by forthcoming solutions. For instance, an attacker could try to use the multihomed device as a means to access another network that would not be normally reachable through the Internet. Even when forwarding to another network is turned off by configuration, an attacker could compromise a system to enable it.

8. Acknowledgments

The authors would like to thank people who have given valuable comments on various multihoming issues on the mailing list, and also those who have suggested directions in the 56th - 61st IETF Meetings. The initial evaluation of NEMO Basic Support on multihoming configurations is a contribution from Julien Charbon.

9. References

9.1. Normative References

- [1] Ernst, T., "Network Mobility Support Goals and Requirements", RFC 4886, July 2007.
- [2] Manner, J. and M. Kojo, "Mobility Related Terminology", RFC 3753, June 2004.
- [3] Ernst, T. and H-Y. Lach, "Network Mobility Support Terminology", RFC 4885, July 2007.
- [4] Devarapalli, V., Wakikawa, R., Petrescu, A., and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", RFC 3963, January 2005.
- [5] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

9.2. Informative References

- [6] Ernst, T., Montavont, N., Wakikawa, R., Ng, C., and K. Kuladinithi, "Motivations and Scenarios for Using Multiple Interfaces and Global Addresses", Work in Progress, October 2006.
- [7] Montavont, N., Wakikawa, R., Ernst, T., Ng, C., and K. Kuladinithi, "Analysis of Multihoming in Mobile IPv6", Work in Progress, February 2006.
- [8] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [9] Arkko, J. and I. Beijnum, "Failure Detection and Locator Pair Exploration Protocol for IPv6 Multihoming", Work in Progress, December 2006.

- [10] Krishnan, S., Montavont, N., Yegin, A., Veerepalli, S., and A. Yegin, "Link-layer Event Notifications for Detecting Network Attachments", Work in Progress, November 2006.
- [11] Narayanan, S., "Detecting Network Attachment in IPv6 Networks (DNav6)", Work in Progress, October 2006.
- [12] Narten, T., Nordmark, E., and W. Simpson, "Neighbor Discovery for IP Version 6 (IPv6)", RFC 2461, December 1998.
- [13] Nordmark, E. and M. Bagnulo, "Level 3 multihoming shim protocol", Work in Progress, November 2006.
- [14] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, May 2000.
- [15] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.
- [16] Huitema, C. and M. Marcelo, "Ingress filtering compatibility for IPv6 multihomed sites", Work in Progress, October 2006.
- [17] Wakikawa, R., Devarapalli, V., and P. Thubert, "Inter Home Agents Protocol (HAHA)", Work in Progress, February 2004.
- [18] Koh, B., Ng, C., and J. Hirano, "Dynamic Inter Home Agent Protocol", Work in Progress, July 2004.
- [19] Tsukada, M., "Analysis of Multiple Mobile Routers Cooperation", Work in Progress, October 2005.
- [20] Miyakawa, S. and R. Droms, "Requirements for IPv6 Prefix Delegation", RFC 3769, June 2004.
- [21] Droms, R. and P. Thubert, "DHCPv6 Prefix Delegation for NEMO", Work in Progress, September 2006.
- [22] Thubert, P. and T.J. Kniveton, "Mobile Network Prefix Delegation", Work in Progress, November 2006.
- [23] Wakikawa, R., "Multiple Care-of Addresses Registration", Work in Progress, June 2006.
- [24] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", RFC 3484, February 2003.

- [25] Thubert, P., Bontous, C., and N. Nicolas, "Nested Nemo Tree Discovery", Work in Progress, November 2006.
- [26] Kumazawa, M., "Token based Duplicate Network Detection for split mobile network (Token based DND)", Work in Progress, July 2005.
- [27] Soliman, H., "Flow Bindings in Mobile IPv6 and NEMO Basic Support", Work in Progress, March 2007.

Appendix A. Alternative Classifications Approach

A.1. Ownership-Oriented Approach

An alternative approach to classifying a multihomed mobile network was proposed by Erik Nordmark (Sun Microsystems) by breaking the classification of multihomed network based on ownership. This is more of a tree-like, top-down classification. Starting from the control and ownership of the HA(s) and MR(s), there are two different possibilities: either (i) the HA(s) and MR(s) are controlled by a single entity, or (ii) the HA(s) and MR(s) are controlled by separate entities. We called the first possibility the 'ISP Model', and the second the 'Subscriber/Provider Model'.

A.1.1. ISP Model

The case of the HA(s) and MR(s) are controlled by the same entity can be best illustrated as an Internet Service Provider (ISP) installing MRs on trains, ships, or planes. It is up to the ISP to deploy a certain configuration of mobile network; all 8 configurations, as described in the Configuration-Oriented Approach, are possible. In the remaining portion of this document, when specifically referring to a mobile network configuration that is controlled by a single entity, we will add an 'ISP' prefix; for example, ISP-(1,1,1) or ISP-(1,n,n).

When the HA(s) and MR(s) are controlled by a single entity (such as an ISP), the ISP can decide whether it wants to assign one or multiple MNPs to the mobile network just like it can make the same decision for any other link in its network (wired or otherwise). In any case, the ISP will make the routing between the mobile networks and its core routers (such as the HAs) work. This includes not introducing any aggregation between the HAs, which will filter out routing announcements for the MNP(s).

To such ends, the ISP has various means and mechanisms. For one, the ISP can run its Interior Gateway Protocol (IGP) over bi-directional tunnels between the MR(s) and HA(s). Alternatively, static routes may be used with the tunnels. When static routes are used, a mechanism to test "tunnel liveness" might be necessary to avoid maintaining stale routes. Such "tunnel liveness" may be tested by sending heartbeats signals from MR(s) to the HA(s). A possibility is to simulate heartbeats using Binding Updates messages by controlling the "Lifetime" field of the Binding Acknowledgment message to force the MR to send Binding Update messages at regular intervals. However, a more appropriate tool might be the Binding Refresh Request

message, though conformance to the Binding Refresh Request message may be less strictly enforced in implementations since it serves a somewhat secondary role when compared to Binding Update messages.

A.1.2. Subscriber/Provider Model

The case of the HA(s) and MR(s) controlled by the separate entities can be best illustrated with a subscriber/provider model, where the MRs belongs to a single subscriber and subscribes to one or more ISPs for HA services. There is two sub-categories in this case: when the subscriber subscribes to a single ISP, and when the subscriber subscribes to multiple ISPs. In the remaining portion of this document, when specifically referring to a mobile network configuration that is in the subscriber/provider model where the subscriber subscribes to only one ISP, we will add an 'S/P' prefix; for example, S/P-(1,1,1) or S/P-(1,n,n). When specifically referring to a mobile network configuration that is in the subscriber/provider model where the subscriber subscribes to multiple ISPs, we will add an 'S/mP' prefix; for example, S/mP-(1,1,1) or S/mP-(1,n,n).

Not all 8 configurations are likely to be deployed for the S/P and S/mP models. For instance, it is unlikely to foresee a S/mP-(*,1,1) mobile network where there is only a single HA. For the S/P model, the following configurations are likely to be deployed:

- o S/P-(1,1,1): Single Provider, Single MR, Single HA, Single MNP
- o S/P-(1,1,n): Single Provider, Single MR, Single HA, Multiple MNPs
- o S/P-(1,n,1): Single Provider, Single MR, Multiple HAs, Single MNP
- o S/P-(1,n,n): Single Provider, Single MR, Multiple HAs, Multiple MNPs
- o S/P-(n,n,1): Single Provider, Multiple MRs, Single HA, Single MNP
- o S/P-(n,1,n): Single Provider, Multiple MRs, Single HA, Multiple MNPs
- o S/P-(n,n,1): Single Provider, Multiple MRs, Multiple HAs, Single MNP
- o S/P-(n,n,n): Single Provider, Multiple MRs, Multiple HAs, Multiple MNPs

For the S/mP model, the following configurations are likely to be deployed:

- o S/mP-(1,n,1): Multiple Providers, Single MR, Multiple HAs, Single MNP
- o S/mP-(1,n,n): Multiple Providers, Single MR, Multiple HAs, Multiple MNPs
- o S/mP-(n,n,n): Multiple Providers, Multiple MRs, Multiple HAs, Multiple MNPs

When the HA(s) and MR(s) are controlled by different entities, it is more likely that the MR is controlled by one entity (i.e., the subscriber), and the MR is establishing multiple bi-directional tunnels to one or more HA(s) provided by one or more ISP(s). In such cases, it is unlikely that the ISP will run IGP over the bi-directional tunnel, since the ISP will most certainly wish to retain full control of its routing domain.

A.2. Problem-Oriented Approach

A third approach was proposed by Pascal Thubert (Cisco Systems). This focused on the problems of multihomed mobile networks rather than the configuration or ownership. With this approach, there is a set of 4 categories based on two orthogonal parameters: the number of HAs, and the number of MNPs advertised. Since the two parameters are orthogonal, the categories are not mutually exclusive. The four categories are:

- o Tarzan: Single HA for Different CoAs of Same MNP

This is the case where one MR registers different CoAs to the same HA for the same subnet prefix. This is equivalent to the case of $y=1$, i.e., the (1,1,*) mobile network.

- o JetSet: Multiple HAs for Different CoAs of Same MNP

This is the case where the MR registers different CoAs to different HAs for the same subnet prefix. This is equivalent to the case of $y=n$, i.e., the (1,n,*) mobile network.

- o Shinkansen: Single MNP Advertised by MR(s)

This is the case where one MNP is announced by different MRs. This is equivalent to the case of $x=n$ and $z=1$, i.e., the (n,*,1) mobile network.

- o DoubleBed: Multiple MNPs Advertised by MR(s)

This is the case where more than one MNPs are announced by the different MRs. This is equivalent to the case of $x=n$ and $z=n$, i.e., the $(n,*,n)$ mobile network.

Appendix B. Nested Tunneling for Fault Tolerance

In order to utilize the additional robustness provided by multihoming, MRs that employ bi-directional tunneling with their HAS should dynamically change their tunnel exit points depending on the link status. For instance, if an MR detects that one of its egress interface is down, it should detect if alternate routes to the global Internet exists. This alternate route may be provided by any other MRs connected to one of its ingress interfaces that has an independent route to the global Internet, or by another active egress interface the MR itself possesses. If such an alternate route exists, the MR should re-establish the bi-directional tunnel using this alternate route.

In the remaining part of this Appendix, we will attempt to investigate methods of performing such re-establishment of bi-directional tunnels. This method of tunnel re-establishment is particularly useful for the $(*,n,n)$ NEMO configuration. The method described is by no means complete and merely serves as a suggestion on how to approach the problem. It is also not the objective to specify a new protocol specifically tailored to provide this form of re-establishments. Instead, we will limit ourselves to currently available mechanisms specified in Mobile IPv6 [5] and Neighbor Discovery in IPv6 [12].

B.1. Detecting Presence of Alternate Routes

To actively utilize the robustness provided by multihoming, an MR must first be capable of detecting alternate routes. This can be manually configured into the MR by the administrators if the configuration of the mobile network is relatively static. It is however highly desirable for MRs to be able to discover alternate routes automatically for greater flexibility.

The case where an MR possesses multiple egress interface (bound to the same HA or otherwise) should be trivial, since the MR should be able to "realize" it has multiple routes to the global Internet.

In the case where multiple MRs are on the mobile network, each MR has to detect the presence of other MR. An MR can do so by listening for Router Advertisement message on its **ingress** interfaces. When an MR receives a Router Advertisement message with a non-zero Router

Lifetime field from one of its ingress interfaces, it knows that another MR that can provide an alternate route to the global Internet is present in the mobile network.

B.2. Re-Establishment of Bi-Directional Tunnels

When an MR detects that the link by which its current bi-directional tunnel with its HA is using is down, it needs to re-establish the bi-directional tunnel using an alternate route detected. We consider two separate cases here: firstly, the alternate route is provided by another egress interface that belongs to the MR; secondly, the alternate route is provided by another MR connected to the mobile network. We refer to the former case as an alternate route provided by an alternate egress interface, and the latter case as an alternate route provided by an alternate MR.

B.2.1. Using Alternate Egress Interface

When an egress interface of an MR loses the connection to the global Internet, the MR can make use of its alternate egress interface should it possess multiple egress interfaces. The most direct way to do so is for the MR to send a binding update to the HA of the failed interface using the CoA assigned to the alternate interface in order to re-establish the bi-directional tunneling using the CoA on the alternate egress interface. After a successful binding update, the MR encapsulates outgoing packets through the bi-directional tunnel using the alternate egress interface.

The idea is to use the global address (most likely a CoA) assigned to an alternate egress interface as the new (back-up) CoA of the MR to re-establish the bi-directional tunneling with its HA.

B.2.2. Using Alternate Mobile Router

When the MR loses a connection to the global Internet, the MR can utilize a route provided by an alternate MR (if one exists) to re-establish the bi-directional tunnel with its HA. First, the MR has to obtain a CoA from the alternate MR (i.e., attach itself to the alternate MR). Next, it sends binding update to its HA using the CoA obtained from the alternate MR. From then on, the MR can encapsulate outgoing packets through the bi-directional tunnel via the alternate MR.

The idea is to obtain a CoA from the alternate MR and use this as the new (back-up) CoA of the MR to re-establish the bi-directional tunneling with its HA.

Note that every packet sent between MNNs and their correspondent nodes will experience two levels of encapsulation. The first level of tunneling occurs between an MR that the MNN uses as its default router and the MR's HA. The second level of tunneling occurs between the alternate MR and its HA.

B.3. To Avoid Tunneling Loop

The method of re-establishing the bi-directional tunnel as described in Appendix B.2 may lead to infinite loops of tunneling. This happens when two MRs on a mobile network lose connection to the global Internet at the same time and each MR tries to re-establish bi-directional tunnel using the other MR. We refer to this phenomenon as tunneling loop.

One approach to avoid tunneling loop is for an MR that has lost connection to the global Internet to insert an option into the Router Advertisement message it broadcasts periodically. This option serves to notify other MRs on the link that the sender no longer provides global connection. Note that setting a zero Router Lifetime field will not work well since it will cause MNNs that are attached to the MR to stop using the MR as their default router too (in which case, things are back to square one).

B.4. Points of Considerations

This method of using tunnel re-establishments is by no means a complete solution. There are still points to consider in order to develop it into a fully functional solution. For instance, in Appendix B.1, it was suggested that MR detects the presence of other MRs using Router Advertisements. However, Router Advertisements are link scoped, so when there is more than one link, some information may be lost. For instance, suppose a case where there are three MRs and three different prefixes and each MR is in a different link with regular routers in between. Suppose now that only a single MR is working; how do the other MRs identify which prefix they have to use to configure the new CoA? In this case, there are three prefixes being announced, and an MR whose link has failed knows that its prefix is not to be used, but it does not have enough information to decide which one of the other two prefixes to use to configure the new CoA. In such cases, a mechanism is needed to allow an MR to withdraw its own prefix when it discovers that its link is no longer working.

Authors' Addresses

Chan-Wah Ng
Panasonic Singapore Laboratories Pte Ltd
Blk 1022 Tai Seng Ave #06-3530
Tai Seng Industrial Estate
Singapore 534415
SG

Phone: +65 65505420
EMail: chanwah.ng@sg.panasonic.com

Thierry Ernst
INRIA
INRIA Rocquencourt
Domaine de Voluceau B.P. 105
Le Chesnay 78153
France

Phone: +33-1-39-63-59-30
Fax: +33-1-39-63-54-91
EMail: thierry.ernst@inria.fr
URI: <http://www.nautilus6.org/~thierry>

Eun Kyoung Paik
KT
KT Research Center
17 Woomyeon-dong, Seocho-gu
Seoul 137-792
Korea

Phone: +82-2-526-5233
Fax: +82-2-526-5200
EMail: euna@kt.co.kr
URI: <http://mmlab.snu.ac.kr/~eun/>

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

Phone: +34 91624 8837
EMail: marcelo@it.uc3m.es

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

