C. Ng
Panasonic Singapore Labs
J. Charbon
Keio and Louis Pasteur University
E. Paik
Seoul National University
T. Ernst
WIDE at Keio University
February 16, 2004

Analysis of Multihoming in Network Mobility Support
draft-ng-nemo-multihoming-issues-03

Status of this Memo

Copyright Notice

Abstract

This document is an analysis of multihoming in the context of network mobility (NEMO). As there are many situations in which mobile networks may be multihomed, we outline possible approaches to classify the multihomed mobile networks. We also describe possible deployment scenarios and we attempt to identify issues that arise when mobile networks are multihomed while mobility supports is taken care by NEMO Basic Support.

Table of Contents

1. Introduction

The goals and objectives of Network Mobility Support (NEMO) are identified in [2] while the terminology is being described in [5]. A solution to provide continuous Internet connectivity to nodes in a mobile network, that is, a network which changes its point of attachment to the Internet, is currently designed by the NEMO Working Group [1]. This solutions basically solves the problem by setting up bi-directional tunnels betweens the mobile routers (MRs) and their Home Agent (HAs), much how this is done in Mobile IPv6 [3], the solution for host mobility.

The purpose of this memo to investigate issues related to such a bi-directional tunneling mechanism when mobile networks are multihomed, i.e. when there is more than one point of attachment between the mobile network and the Internet (see definitions in draft [1]). Goals and objectives of multihoming are discussed in a separate document [6] with fits to both fixed nodes, mobile nodes, fixed networks and mobile networks. Our objectives are three-folds:

- o To capture issues for deploying a multihomed mobile network
- o To identify which multihoming configurations are useful
- o To identify issues in NEMO Basic Support that prevent to support the useful configurations. It doesn't mean that those not supported will not work with NEMO Basic Support, just that it is up to the implementors to make it work (hopefully issues discussed

in the document will be helpful to these implementors).

For doing so, Section 2 first outlined several taxonomies to classify multihomed mobile networks. This section outlines 3 different approaches to classifying multihomed mobile network. Benefits and issues of multihoming peculiar to network mobility support are discussed in Section 3. Next, we described deployment scenarios of multihomed mobile networks in Section 4. Following this, we study the general issues, and we conclude with an evaluation of NEMO Basic Support for multihomed configurations.

In order to understand this memo, the reader is expected to be familiar with the aboved cited documents, i.e. with the NEMO terminology [5], Goals and Objectives of Multihoming [6], Goals and Requirements of Network Mobility Support [2], and the NEMO Basic Support specification [1].

2. Classification

Various discussions on the topic of multihoming issues in NEMO has been carried out on the Mailing List. As there are several configurations in which mobile networks are multihomed, there is a need to classify multihomed mobile network into a clearly defined taxonomy. This can be done in various ways. Three main approaches have been proposed on the NEMO mailing list. These are, namely, (i) Configuration-Oriented Approach, (iii) Ownership-Oriented Approach, and (ii) Problem-Oriented Approach. As the WG consensus seems to have converged to the Configuration-Oriented dApproach, we described only this approach here. The other two appraoches can be found in Appendix A.1 and Appendix A.2.

Configuration-Oriented Approach

Multihomed configurations can be classified depending on how many mobile routers are present, how many egress interfaces and home addresses the mobile routers have, how many prefixes (NEMO-prefixes) are advertised to the mobile network nodes, etc. For doing so, we use three key parameters differentiating different multihomed

configurations. With these parameters, we can refer to each configuration 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 a mobile network has only a single mobile router. presumably with multiple egress interfaces or multiple home addresses.

x=N implies a mobile network has more than one mobile router advertising an egress route.

o 'y' indicates the number of HAs associated with the mobile network, where:

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

y=N implies that more than one home agents (possibly in different administrative domains) are assigned the mobile network.

o 'z' indicates the number of NEMO-prefix announced to MNNs, where:

z=1 implies that a single NEMO-prefix is advertised to the mobile network nodes.

z=N implies that more than one NEMO-prefix are advertised to the mobile network nodes.

It can be seen that the above three parameters are fairly orthogonal to one another. Thus different values of 'x', 'y' and 'z' give rise to different combinations of the 3-tuple (x,y,z). As described below, a total of 8 possible configurations can be identified.

2.1 (1,1,1): Single MR, Single HA, Single Prefix

The (1,1,1) mobile network has only one mobile router advertising a single NEMO-prefix. In addition, the mobile router associates with only one home agent at any one time. This makes the mobile network

very similar to a non-multihomed mobile network, except for the fact that the mobile router may either (i) use more than one egress links at the same time, or (ii) use more than one home address at the same time.

Since only one NEMO-prefix is advertised, the mobile network nodes are (usually) not multihomed.

```
                            -----
            _    p      _  |     |
           |_|-|<-_  |-|_|-|     |-|        _
            _  |-|_|=|     |_____| |  _  |-|_|
           |_|-|     |             |-|_|-|
                                         |
           MNNs   MR   AR  Internet  AR    HA
```

Figure 2.1 - (1,1,1) Multihomed Mobile Network

2.2 (1,1,N): Single MR, Single HA, Multiple Prefixes

The (1,1,N) mobile network has only one mobile router, which associates to only one home agent at any one time. However, two or more NEMO-prefixes are advertised to the mobile network nodes. No associations is assumed between the NEMO-prefixes and the home addresses of the mobile router.

Since a plurality of NEMO-prefixes are advertised, mobile network nodes can generally be multihomed themselves, where each mobile network node is allocated one address in each NEMO-prefix.


```
                            -----
            _   p1,p2   _  |     |
           |_|-|<-_  |-|_|-|     |-|        _
            _  |-|_|=|     |_____| |  _  |-|_|
           |_|-|     |             |-|_|-|
                                         |
           MNNs   MR    AR Internet  AR    HA
```

Figure 2.2 - (1,1,N) Multihomed Mobile Network

2.3 (1,N,1): Single MR, Multiple HAs, Single Prefix

The (1,N,1) mobile network has only one mobile router advertising a single NEMO-prefix. The mobile router, however, associates to multiple home agents, possibly one home agent per home addresses. No assumption is made on whether or not the home agents belongs to the same administrative domain.

Since only one NEMO-prefix is advertised, the mobile network nodes are (usually) not multihomed.

```
                                        AR    HA2
                                      _  |
                                   |-|_|-|  _
                            -----  |     |-|_|
            _    p      _  |     |-|
           |_|-|<-_  |-|_|-|     |
            _  |-|_|=|     |_____|-|        _
           |_|-|     |             |  _  |-|_|
                                   |-|_|-|
                                         |
           MNNs  MR    AR  Internet  AR    HA1
```

Figure 2.3 - (1,N,1) Multihomed Mobile Network

2.4 (1,N,N): Single MR, Multiple HAs, Multiple Prefixes

The (1,n,n) mobile network has only one mobile router. However, the mobile router advertises more than one NEMO-prefix, and also associates to multiple home agents at the same time, possibly one home agent per home address. No assumptions is made on whether or not the home agents belongs to the same administrative domain.

Since a plurality of NEMO-prefixes are advertised, mobile network nodes can generally be multihomed themselves, where each mobile network node is allocated one address in each NEMO-prefix.

```
                                     AR    HA2
                                      _  |  _
                            _____  |-|_|-|-|_|
            _   p1,p2   _  |     |-|     |
           |_|-|<-_  |-|_|-|     |          _
            _  |-|_|=|     |_____|-|  _  |-|_|
           |_|-|     |             |-|_|-|
                                   |     |
           MNNs  MR    AR  Internet  AR    HA1
```

Figure 2.4 - (1,N,N) Multihomed Mobile Network

2.5 (N,1,1): Multiple MRs, Single HA, Single Prefix

The (N,1,1) mobile network has more than one mobile router advertising global routes. These mobile routers, however, advertise the same NEMO-prefix and associate to the same home agent. Since only one NEMO-prefix is advertised, the mobile network nodes are (usually) not multihomed.

```
                 MR2
               p<-_  |
            _  |-|_|-|  _____
           |_|-|     |-|     |
            _  |       |     |-|        _
           |_|-|  _  |-|_____| |  _  |-|_|
               |-|_|-|         |-|_|-|
               p<-   |               |
           MNNs  MR1   Internet  AR    HA
```

Figure 2.5 - (N,1,1) Multihomed Mobile Network

2.6 (N,1,N): Multiple MRs, Single HA, Multiple Prefixes

The (N,1,N) mobile network has more than one mobile router advertising different global routes and different NEMO-prefixes. However, these mobile routers associate to the same home agents.

Since a plurality of NEMO-prefixes are advertised, mobile network nodes can generally be multihomed themselves, where each mobile network node is allocated one address in each NEMO-prefix.

```
                 MR2
              p2<-_  |
            _  |-|_|-|  _____
           |_|-|     |-|     |
            _  |       |     |-|        _
           |_|-|  _  |-|_____| |  _  |-|_|
               |-|_|-|         |-|_|-|
              p1<-   |               |
           MNNs  MR1   Internet  AR    HA
```

       Figure 2.6 - (N,1,N) Multihomed Mobile Network

2.7 (N,N,1): Multiple MRs, Multiple HAs, Single Prefix

   The (N,N,1) mobile network has more than one mobile router advertising different global routes. The mobile routers are also associated to more than one home agents at any one time. No assumptions is made on whether or not the home agents belongs to the same administrative domain. However, the mobile routers advertises the same NEMO-prefix. Since only one NEMO-prefix is advertised, the mobile network nodes are (usually) not multihomed.

```
                 MR2             AR    HA2
                 p                _  |
                <-_  |         |-|_|-|  _
            _  |-|_|-|  _____  |     |-|_|
           |_|-|     |-|     |-|
            _  |       |     |
           |_|-|  _  |-|_____|-|        _
               |-|_|-|         |  _  |-|_|
                <-   |         |-|_|-|
                 p                   |
           MNNs  MR1   Internet  AR    HA1
```

       Figure 2.7 - (N,N,1) Multihomed Mobile Network

2.8 (N,N,N): Multiple MRs, Multiple HAs, Multiple Prefixes

   The (N,N,N) mobile network has more than one mobile router advertising different global routes and different NEMO-prefixes. The mobile routers are also associated to more than one home agent at any

one time. No assumptions is made on whether or not the home agents belongs to the same administrative domain.

Since a plurality of NEMO-prefixes are advertised, mobile network nodes can generally be multihomed themselves, where each mobile

network node is allocated one address in each NEMO-prefix.

```
                 MR2             AR    HA2
                 p2               _  |
                <-_  |         |-|_|-|  _
            _  |-|_|-|  _____  |     |-|_|
           |_|-|     |-|     |-|
            _  |       |     |
           |_|-|  _  |-|_____|-|        _
               |-|_|-|         |  _  |-|_|
                <-   |         |-|_|-|
                 p1                  |
           MNNs  MR1   Internet  AR    HA1
```

Figure 2.8 - (N,N,N) Multihomed Mobile Network

3. Benefits/Issues of Multihoming in NEMO

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

1. Ubiquitous Acccess

2. Redundancy/Fault-Recovery

3. Load Sharing

4. Load Balancing

5. Ubiquity

6. Preference Settings

This section discusses these from a NEMO perspective and we give typical instances for each case of our taxonomy.

Mobile networks are typically connected by means of wireless and thus less reliable links. In addition, there could be many nodes behind the MR, so a failure to connect to the Internet has a more important impact than once only one node is concerned by a lack a failure or loss of connectivity. Real-life scenarios highlighted in [6] have illustrated that offering a permanent access to mobile networks such as vehicles typically require the use of several interfaces and technologies since the mobile network may be moving in distant geographical locations where different access technologies are

provided and gouverned by distinct access control policies.

3.1 Deployment Scenarios

Here, we list some example scenarios for each configurations

x=1: Multihomed mobile network with one mobile router

- o A mobile router with dual/multiple access interfaces (i.e. 802.11 and GPRS capabilities). When it subscribed to same ISP for both accesses, this is a S/P-(1,1,*). If it different ISPs are offering the two accesses independently, this is a S/mP-(1,N,N).

  Benefits: Ubiquity, Redundancy/Fault-Recovery

x=N: Mutlihomed mobile networks with multiple mobile routers

- o A train with one MR in each car. This is usually served by same home agent, thus usually a (N,1,*). Alternatively, the train company might be forced to use different ISP when the train go to different locations, thus it is a (N,N,N).

  Benefits: Load Sharing

- o A W-PAN with a GPRS_enabled phone and a WiFi-enabled PDA. The two access technology are usually separately subsribed, thus it is likely to be S/mP-(N,N,N).

  Benefits: Ubiquity, Redundancy/Fault-Recovery

y=1: Multihomed mobile networks with one home agent

- o Most single ISP cases in above examples.

y=N: Multihomed mobile networks with multiple home agents

o Most multiple ISP cases in above examples.

o A transalantic flight that cahnge its home agent when its in different continents. This is a (1,N,1) network if there is only one mobile router.

Benefits: Ubiquity (reduce delays)

z=1: Multihomed mobile networks with one prefix

o Most single ISP cases in above examples.

z=N: Multihomed mobile networks with multiple prefixes

o Most multiple ISP cases in above examples.

o A car with a prefix taken from my home (I pay the traffic on thisprefix) and one that belong to the car-manufacturer (for maintenance, traffic is paid by the car-manufacturer [indeed me when I buy the car:-)]). This will typically be a (1,1,N).

Benefits: preference settings

4. Problem Statement

4.1 Connection Availability

Multiple connections of MRs can be used simultaneously or one at a time.

o When multiple connections are used simultaneously, the mode of operation can be either primary-secondary or peer-to-peer. These configurations can be useful especially for large mobile networks, but there are many implementation issues which need to be addressed, e.g. which connection will be selected for each traffic folw that goes into/out of the mobile network ?

o When only one connection can be used at a time, e.g. in the case

where a single connection has to substitute for all of the other failed connections, a connection selection mechanism is needed. The connection selection can depend on which connection is available at that time.

4.2 Connection Selection

The connection can be selected by the home agent (HA), the MR, and/or the MNN.

- o The HA can select a connection based on the binding update information in the binding cache.
- o The MR can select a connection since the MR is one of the main bodies of the connection.
- o The MNN should be able to select a connection, e.g. in case where a user wants to select a particular access technology among the available technologies for reasons of cost or data rate.
- o A hybrid mechanism should be also available, e.g. one in which the HA, the MR, and/or the MNN coordinate to select a connection.

4.3 Scalability

Should a new solution meets the all the eight configurations and the scenarios mentioned in section 3 ?

4.4 Route Optimization Considerations

RO problems in multihomed mobile networks are dependant on how the

connections are available and selected.

- o In case of multiple HAs and HoAs, new route optimization may be possible by routing between CN and multiple HAs with different HoAs.
- o When multiple connections are available simultaneously, how the CN knows about the availability and optimizes route ?

4.5 Ingress Filtering

To enjoy the benefits of multihoming as descriebd earlier, it is often necessary to divert packets from the same session between two different bi-directional tunnels. This is especially treu when we consider the fault recovery feature of multihoming when packets froma failed bi-directional tunnel is sent via an alternative (perhaps newly established) bi-directional tunnel.

> When doing so, care has to be taken to prevent ingress filtering from dropping the outgoing packets when the two tunnels end at different home agents. Ingress filtering occurs when different mobile network prefixes are handled by different home agents. For example, consider the case when a mobile network has two tunnel connections to home agents HA1 and HA2. The mobile network prefix P1 is registered to HA1, and mobile network prefix P2 is registered to HA2. Mobile network nodes are free to auto-configure their addresses based on any of P1 or P2. When the tunnel to HA1 is broken, packets sent through the tunnel to HA1 are diverted to send through the tunnel to HA2. If HA2 (or some border gateway in the domain of HA2) performs ingress filtering, packets will a source address prefix of P1 may be discarded.

To avoid ingress filtering for such cases, the mobile router(s) can stop advertising the network prefix P1. This will stop mobile network node from using source address auto-configured from prefix P1. However, such a method suffers from the following two limitations:

o Switching of source address is a long process since nodes have to wait for source address to get deprecated [7].

o In addition, switching of source address will force transport sessions without multihoming capabilities (such as TCP) to be terminated, and re-established using the new source address. Transport sessions with multihoming capabilities (such as SCTP) may be able to continue without disruption.

It is possible to overcome these limitations by using nested tunnels.

Appendix B describes one such approach.

4.6 Failure Detection

In order for fault recovery to work, the mobile routers and home agents must first possess a means to detect failures. It is expected for faults to occur more readily at the end of the mobile network, due to the use of wireless connections. The mobile router can then rely on router advertisements from access routers, or other layer two trigger mechanisms to detect faults. In comparison, it is more difficult for home agents to detect tunnel failures. For an ISP deployment model, the home agents and mobile routers can use proprietary methods (such as constant transmission of heartbeat signals) to detect failures and check tunnel liveness. In the S/P model, a lack of standardized "tunnel liveness" protocol means that it is harder to detect failures.

A possible method is for the mobile routers to send binding updates more regularly with shorter Lifetime value. Similarly the home agents can return binding acknowledgment messages with smaller Lifetime values as well, thus forcing the mobile routers 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 home agents must be encrypted.

5. Evaluation of Basic NEMO Solution

This section, we attempt to analyze what are the problems faced in each of the 8 categories. It shouldn't matter if some of the categories share the same problem(s).

- (1,1,1) Mobile Network

  The (1,1,1) mobile network has only one mobile router registering more than one care-of-addresses to the same home agent, and advertising only one prefix. The mobile router can either have more than one care-of-addresses bound to the same home-address, or it can have various care-of-address and home-address pairs.

  Either way, this is a MIPv6 problem. Multiple pairs of different care-of-address and home-address is perfectly alright with MIPv6. The fact that they specify the same NEMO-prefix in binding updates shouldn't cause a problem either. Having a home-address bound to multiple care-of-address simultaneously may be a problem for MIPv6. This will require a solution like [8].

- (1,1,N) Mobile Network

  The (1,1,N) mobile network is similar to the (1,1,1) mobile network, and thus face the same problem when there is only one home-address bound to multiple care-of-addresses. In addition, it is possible for the MR to include multiple mobile network prefix options in a single binding update, thus having multiple network prefixes should not create additional issues.

- (1,N,1) Mobile Network

  The (1,N,1) mobile network has one mobile router registering to multiple home agents. There is the question of whether a mobile router can register the same home-address to different home agents simultaneously with the 'H' bit set. If not, the mobile router can only register different home-address and care-of-address pairs to different home agents. In any case, this is a MIPv6 issue.

  The NEMO-specific problem is the fact that a NEMO-prefix has a care-of in different home agents. It might be possible that only one home-agent will actively advertise a route to the NEMO-prefix. The case of multiple home agents at different domains advertising a route to the same NEMO-prefix may pose a problem in the routing infrastructure as a whole. The implications of this aspect needs further exploration.

o (1,N,N) Mobile Network

The (1,N,N) mobile network has one mobile router registering to multiple home agents multiple NEMO-prefixes. The same question of whether the same home-address can be simultaneously registered to multiple home agents.

This (1,N,N) network can avoid the problem of registering care-ofs for the same prefix to different home agents by registering care-of for one prefix at one home-agent.

o (N,1,1) Mobile Network

The (N,1,1) mobile network has two or more active egress mobile routers, registering to same home agents, and advertising the same prefix. May not have any problem at all if the mobile routers are manually configured to announce the same prefix. It is also possible that prefix delegation is used to ensure all routers advertise the same NEMO-prefix since all routers are handled by the same home agent. The home-agent will see two HoA-CoA pairs taking care of the same NEMO-prefix.

o (N,1,N) Mobile Network

The (N,1,N) mobile network has multiple active egress mobile routers registering to the same home-agent, and advertising multiple prefixes. If a mobile router is advertising more than one prefix, we have the same problem as (1,1,N) as in how to register more than one NEMO-prefix to the same home-agent.

On the other hand, if each mobile router take cares of a separate (and only one) NEMO-prefix, then there should not be any NEMO-specific problem.

o (N,N,1) Mobile Network

The (N,N,1) mobile network has multiple mobile routers registering to different home agents, but advertising the same prefix. There is the same issues as in (1,N,1) of a NEMO-prefix having a care-of in different home agents. In addition, there is a question how to perform prefix delegation such that two home agents will delegate the same prefix to different mobile routers. Certain level of home-agent co-ordination may be required here.

o (N,N,N) Mobile Network

The (N,N,N) mobile network has multiple mobile routers, registering to multiple home-agents and advertising prefixes. This may be a case of multiple non-multihomed network superimposed together, i.e. each mobile router take cares of one prefix, and register to separate home agents.

On the other hand, if one mobile router takes cares of more than one prefix, we have similar problems as (1,1,N) and (N,1,N). In addition, if more than one mobile router takes care of the same prefix, we have similar issues as (N,N,1). In any case, we see that the problems within this configurations can be decomposed into problems from other configurations.

6. 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 - 58th IETF Meetings.

References

[1] Devarapalli, V., "Nemo Basic Support Protocol", draft-ietf-nemo-basic-support-02 (work in progress), December 2003.

[2] Ernst, T., "Network Mobility Support Goals and Requirements", draft-ietf-nemo-requirements-02 (work in progress), Feb 2004.

[3] Johnson, D., Perkins, C. and J. Arkko, "Mobility Support in IPv6", draft-ietf-mobileip-ipv6-24 (work in progress), July 2003.

[4] Simpson, W., "IP in IP Tunneling", IETF RFC 1853, October 1995.

[5] Ernst, T. and H. Lach, "Network Mobility Support Terminology", draft-ietf-nemo-terminology-01 (work in progress), Feb 2004.

[6] Ernst, T., "Goals and Benefits of Multihoming", draft-ernst-generic-multihoming-00 (work in progress), February 2004.

[7] Draves, R., "Default Address Selection for Internet Protocol version 6 (IPv6)", IETF RFC 3484, February 2003.

[8] Wakikawa, R., "Multiple Care-of Addresses Registration", draft-wakikawa-mobileip-multiplecoa-02 (work in progress), September 2003.

[9] Narten, T., Nordmark, E. and Simpson, W., "Neighbour Discovery for IPv6", IETF RFC 2461, December 1998.

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: cwng@psl.com.sg

Julien Charbon
Keio University, Louis Pasteur University
Keio University.
5322 Endo
Fujisawa-shi, Kanagawa 252-8520
JP

Phone: +81-466-49-1100
Fax: +81-466-49-1395
EMail: julien@sfc.wide.ad.jp
URI: http://www.sfc.wide.ad.jp/~julien/

Paik, Eun Kyoung
Seoul National University
Multimedia Communications Lab., Seoul National Univ.
Shillim-dong, Kwanak-gu
Seoul 151-744
Korea

Phone: +82-2-880-1832
Fax: +82-2-872-2045
EMail: eun@mmlab.snu.ac.kr
URI: http://mmlab.snu.ac.kr/~eun/

Ernst Thierry
WIDE at Keio University
Jun Murai Lab., Keio University.
K-square Town Campus, 1488-8 Ogura, Saiwa-Ku
Kawasaki, Kanagawa 212-0054
Japan

Phone: +81-44-580-1600
Fax: +81-44-580-1437
EMail: ernst@sfc.wide.ad.jp
URI: http://www.sfc.wide.ad.jp/~ernst/

Appendix A. Alternative Classifications Approach

A.1 Ownership-Oriented Approach

An alternative approach to classifying multihomed mobile network is proposed by Eric 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 posibility 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 mobile routers 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 network prefixes 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 home agents) work. This include not introducing any aggregation between the home agents which will filter out routing announcements for the mobile prefix(es).

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 interval. 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) are controlled by the separate entities can be best illustrated with a subscriber/provider model, where the mobile routers belongs to a single subscriber and subscribes to one or more ISPs for home agent 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 Prefix
- o S/P-(1,1,N): Single Provider, Single MR, Single HA, Multiple Prefixes
- o S/P-(1,N,1): Single Provider, Single MR, Multiple HAs, Single Prefix
- o S/P-(1,N,N): Single Provider, Single MR, Multiple HAs, Multiple Prefixes
- o S/P-(N,N,1): Single Provider, Multiple MRs, Single HA, Single Prefix
- o S/P-(N,1,N): Single Provider, Multiple MRs, Single HA, Multiple Prefixes
- o S/P-(N,N,1): Single Provider, Multiple MRs, Multiple HAs, Single Prefix
- o S/P-(N,N,N): Single Provider, Multiple MRs, Multiple HAs, Multiple Prefixes

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 Prefix

- o S/mP-(1,N,N): Multiple Providers, Single MR, Multiple HAs, Multiple Prefixes

- o S/mP-(N,N,N): Multiple Providers, Multiple MRs, Multiple HAs, Multiple Prefixes

When the HA(s) and MR(s) are controlled by different entities, it is more likely the scenario where 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 case, it is unlikely for the ISP to run IGP over the bi-directional tunnel, since ISP would most certainly wish to retain full control of its routing domain.

A.2 Problem-Oriented Approach

A third approach is proposed by Pascal Thubert (Cisco System). 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 home agents, and the number of subnet prefixes advertised. Since the two parameters are orthogonal, the categories are not mutually exclusive. The four categories are:

- o Tarzan: Single HA for Different Care-ofs of Same Prefix

  This is the case where one mobile router registers different care-of-addresses to the same home agent for the same subnet prefix. This is equivalent to the case of y=1, i.e. the (1,1,N) mobile network.

- o JetSet: Multiple HA for Different Care-ofs of Same Prefix

  This is the case where the mobile router registers different care-of-addresses to different home agents for the same subnet prefix. This is equivalent to the case of y=N, i.e. the (1,N,*) mobile network.

- o Shinkansen: Single Prefix Advertised by Mobile Router(s)

  This is the case where one subnet prefix is announced by different mobile routers. This is equivalent to the case of z=N, i.e. the (1,*,N) mobile network.

- o DoubleBed: Multiple Prefixes Advertised by Mobile Router(s)

  This is the case where more than one subnet prefixes are announced by the different mobile routers. This is equivalent to the case of 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 multi-homing, mobile routers that employ bi-directional tunneling with their home agents should dynamically change their tunnel exit points depending on the link status. For instance, if a mobile router detects that one of its egress interface is down, it should detect if any other alternate route to the global Internet exists. This alternate route may be provided by any other mobile routers connected to one of its ingress interfaces that has an independent route to the global Internet, or by another active egress interface the mobile router it self possess. If such an alternate route exists, the mobile router should re-establish the bi-directional tunnel using this alternate route.

In the remaining part of this section, we will attempt to investigate methods of performing such re-establishment of bi-directional tunnels. It is not the objective of this memo 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 and Neighbor Discovery in IPv6 [9].

B.1 Detecting Presence of Alternate Routes

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

The case where a mobile router possesses multiple egress interface (bound to the same home agent or otherwise) should be trivial, since the mobile router should be able to "realize" it has multiple routes to the global Internet.

In the case where multiple mobile routers are on the mobile network, each mobile router has to detect the presence of other mobile router. A mobile router can do so by listening for Router Advertisement message on its *ingress* interfaces. When a mobile router receives a Router Advertisement message with a non-zero Router Lifetime field from one of its ingress interfaces, it knows that another mobile router which can provide an alternate route to the global Internet is present in the mobile network.

B.2 Re-Establishment of Bi-Directional Tunnels

When a mobile router detects that the link be which its current bi-directional tunnel with its home agent 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 mobile router; secondly, the alternate route is provided by another mobile router 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 mobile router.

B.2.1 Using Alternate Egress Interface

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

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

B.2.2 Using Alternate Mobile Router

When the mobile router loses a connection to the global Internet, the mobile router can utilize a route provided by an alternate mobile router (if one exists) to re-establish the bi-directional tunnel with its home agent. First, the mobile router has to obtain a care-of-address from the alternate mobile router (i.e. attaches itself to the alternate mobile router). Next, it sends binding update to its home agent using the care-of-address obtained from the alternate mobile router From then on, the mobile router can encapsulates outgoing packets through the bi-directional tunnel using via the alternate mobile router.

The idea is to obtain a care-of-address from the alternate mobile router and use this as the new (back-up) care-of-address of the mobile router to re-establish the bi-directional tunneling with its home agent.

Note that every packet sent from/to mobile network nodes to/from

their correspondent nodes will experience two levels of encapsulation. First level of tunneling is done between a mobile router which the mobile network node uses as its default router and the mobile router's home agent. The second level of tunneling is done between the alternate mobile router and its home agent.

B.3 To Avoid Tunneling Loop

The method of re-establishing the bi-directional tunnel as described in Section 3.2 may lead to infinite loops of tunneling. This happens when two mobile routers on a mobile network lose connection to the global Internet at the same time and each mobile router tries to re-establish bi-directional tunnel using the other mobile router. We refer to this phenomenon as tunneling loop.

One approach to avoid tunneling loop is for a mobile router 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 mobile routers 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 mobile network nodes that are attached to the mobile router to stop using the mobile

router as an access router too (in which case, things are back to square one).

B.4 Other Considerations

When a mobile network is multihomed, mobile network nodes may receive Router Advertisements that advertise different network prefixes. This is usually the case when the multihomed mobile network has two or more mobile routers advertising different routes to the global Internet. It may also occur when the mobile network has only one mobile router with multiple egress interfaces bound to different home agents. In these situations, mobile network nodes typically only select one to form its global (possibly care-of) address.

In view of this, it may be desirable for mobile network node to be able to acquire "preference" information on each mobile network prefix from the mobile routers. This allows default address selection mechanism such as that specified in [7] 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.

Appendix C. Change Log

- o Changes from version -02 to version -03

    * Merged with draft-eun-nemo-multihoming-problem-statement (see "Problem Statement" (Section 4))

    * Included conclusions from draft-charbon-nemo-multihoming-evaluation

    * Re-organize some part of "Benefits/Issues of Multhoming in NEMO" to "Problem Statement" (Section 4)

    * Remove lot of text to be in sync with [6].

* Title change from "Multihoming Issues in NEMO Basic Support" to "Analysis of Multihoming in NEMO"

* Changed (w,x,y) to (x,y,z) in taxonomy.

* Moved alterntaive approaches of classification to Appendix

* Creation of this Change-Log itself ;-)

Intellectual Property Statement

Full Copyright Statement

Acknowledgment

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