

Multihoming with NEMO Basic Support

Thierry Ernst[†] and Julien Charbon^{†‡}

[†]Keio University, Jun Murai Lab
5322 Endo, Fujisawa-shi, Kanagawa 252-8520, Japan, ernst@sfc.wide.ad.jp

[‡]Louis Pasteur University
7, rue René Descartes, 67084 Strasbourg Cedex, France, julien@sfc.wide.ad.jp

ABSTRACT

The IETF's NEMO working group has been set up recently in an effort to address issues related to network mobility, *i.e.* entire IPv6 networks that change their point of attachment to the Internet topology such as Personal Area Networks (PANs), access networks deployed in public transportation, and networks of sensors embedded in private vehicles. Network mobility is thus the opportunity to realize the ubiquitous Internet, *i.e.* permanent access anywhere at anytime, in fixed locations and while on the move, provided that any available access network can be accommodated. For doing so, mobile networks may be multihomed, *i.e.* having multiple points of attachment to the Internet. This provides for redundancy, load-sharing, and policy-routing. The purpose of this paper is to analyze the behavior of NEMO Basic Support – the IETF solution to manage mobility of networks – when mobile networks are multihomed and to evaluate its ability to support such configurations. Possible multihomed configurations are first classified according to a few parameters. For each case, we identify missing mechanisms and we make a number of recommendations that we proposed to the IETF.

Keywords: IPv6, Network Mobility, Multihoming, NEMO.

1 INTRODUCTION

In order to achieve the ubiquitous Internet, connection to the network must be available everywhere, at anytime, without disruption of service. Internet access at office and home is now common, but we currently lack Internet access in the streets and when commuting. However, with the advent of mobile technologies in IPv6, and particularly network mobility support and multihoming, we have a chance to realize the ubiquitous Internet we are waiting for.

Thanks to network mobility support which allows an entire network, referred to as a *mobile network*, to migrate in the Internet topology, anything can be connected to the Internet. Cases of mobile networks include PANs (small networks attached to people and composed of Internet appliances like PDAs, mobile phones, digital cameras, etc.), networks of sensors deployed in vehicles (aircrafts, boats, buses, trains), and access networks deployed in public transportation (taxis, trains, aircrafts, trucks and personal cars) to provide Internet access to devices carried by their passengers (laptop, camera, mobile phone, and even PANs).

However, to ensure continuous connectivity to the Internet, at anytime, any place, the mobile network is preferably best

connected via several interfaces, several access technologies and to distinct access networks, which we refer to as *multihomed mobile network*. Distinct interfaces may indeed be active simultaneously. As a result, the system must be able to deal with both *horizontal handovers* (between access points using the same communication medium) and *vertical handovers* (between distinct communication medium). This is particularly true from an in-vehicle network standpoint since vehicles will typically be connected to the Internet via multiple wireless medium [14]. As a wide coverage nor a lack of failure cannot be ensured by a single *Internet Service Provider* (ISP), handovers may need to be performed between distinct administrative domains and thus topologically distant parts of the Internet (*global mobility*). For instance, this either may occur when a vehicle crosses country boundaries, or when access is offered by different ISPs. The latter is effective to keep communication costs low from a user's point of view. Having multiple points of attachment is also effective to avoid disruption of service when either a particular technology is not available in the geographic area or when one is experiencing some sort of failure. For instance, one can use 802.11b near parking lots or when traffic jams occur, whereas DSRC (Dedicated Short Range Communication) would be used on highways and the mobile phone in areas with typically low traffic density.

Network mobility support combined with multihoming is thus the missing piece which will allow us to be permanently tuned to the Internet even while on the move between our house and office.

The purpose of the present paper is thus to study multihoming issues related to network mobility. In the following sections, we first overview what is network mobility together with the recent work conducted on this topic at the IETF. Then, we overview what are the benefits of multihoming. Since multihoming covers a number of configurations, these are first classified by means of a taxonomy. We then analyze the IETF solution's ability to support multihoming for each case and we outline potential mechanisms to address the issues. Our recommendations are summarized before concluding with this paper.

2 NETWORK MOBILITY OVERVIEW

Mobility in the Internet arises when a portion of the network changes its point of attachment to the overall topology. However, the Internet is hardly tuned to allow mobility in the midst of data transfers because protocols have not been con-

ceived for devices that change their point of attachment in the topology. Each device is identified by a unique IPv6 address with a prefix which shows the location of the devices in the Internet topology. There is typically a change of this physical IPv6 address each time a *mobile node (MN)* changes its point of attachment and thus its reachability in the Internet topology. Support mechanisms are then necessary to maintain open connections. In some cases, the mobile node may indeed be a *mobile router* with a number of nodes (*MNNs*) attached behind it. This forms what we refer to as a *mobile network*. Only the MR changes its physical point of attachment and thus its address. However, this changes of address has an impact on routing to the entire mobile network. This results in losing packets in transit and breaking transport protocols connections if mobility is not handled by specific services.

Traditional work in mobility support is to provide continuous and uninterrupted Internet access to *mobile hosts*. Host mobility support in IPv6 is performed by Mobile IPv6 [21] and is handled by the Mobile IP Working Group at the IETF (Internet Engineering Task Force). On the other hand, the question of networks that frequently change their point of attachment to the Internet has only gained the deserved attention from the Internet community for about two years, when the topic was initiated in the Mobile IP Working Group. Network Mobility support is now handled by the NEMO Working Group set up at the IETF in October 2003. NEMO's primary objective is to preserve session continuity between CNs and all MNNs behind the MR while the MR changes its point of attachment.




Fig.1: NEMO Terminology

The first task of the working group was to define a terminology [11] which we use in this paper. Basic terms are illustrated on Fig.1. A MR has at least two interfaces, the *egress interface* is first attached to the home link, and later to a visited link. The *ingress interface* is attached to an *internal link* in the mobile network. All nodes (MR and *MNNs*) attached to a given internal link have their addresses taken from the same mobile network prefixes (MNP) advertised on this link. *MNNs* are either *fixed nodes* or *mobile nodes*. Fixed nodes are unable to change their point of attachment while keeping

their connections open, whereas mobile nodes have this ability, presumably using Mobile IPv6. If such a mobile node is indeed a MR with a number of nodes behind it, a sub-mobile network is getting attached to a root-mobile network. In this former case, the aggregated network is said to be *nested*.

The working group is working on a solution termed *NEMO Basic Support* [6] which inherits from earlier proposals, including ours [9], [13], [12]. This protocol associates each egress interface of a MR with two distinct addresses, much like what is done in Mobile IPv6. The *home address (HoA)* serves as a permanent location invariant identifier whereas the *care-of address (CoA)* serves as a routing directive to the current point of attachment. The permanent home address MR_{HoA} is obtained in the home network and has the same prefix as the home link. The temporary care-of address MR_{CoA} is obtained in the visited network and formed based on the prefix advertised on the visited link.

The purpose of the protocol is to establish bi-directional tunnels between the home links and the mobile network for each couple (MR_{HoA}, MR_{CoA}). For doing so, bindings between the MNP and the corresponding MR_{CoA} are registered with the Home Agent (a router located on the home link) by means of a *Binding Update* similar to Mobile IPv6's. A new flag 'R' is indeed added in the Mobility Header (Fig.2) to indicate the Binding Update is coming from a MR (i.e. a Prefix-BU). Actual information about the MNP(s) owned by the MR is either contained in a Mobile Network Prefix Option (Fig.3) (i.e. explicit mode) or in a Mobile Network Prefix Length Option (Fig.4) (i.e. explicit combined mode) or not advertised at all (i.e. implicit mode). As a result of a successful registration, packets originated from a CN and sent to a MNN are transmitted to the home link where they are encapsulated to MR_{CoA} . At this point, the packet is decapsulated by MR and forwarded in the mobile network.

The accustomed reader may have noticed that NEMO Basic Support does not allow direct routing between CNs and MRs. The working group has decided to leave routing optimizations and other challenging issues for later at the IETF because these are considered too complex to be solved quickly. Solutions for routing optimization do exist, but are too immature for standardization. Most documents related to network mobility can be retrieved from the NEMO web page [1].




Fig.2: NEMO Mobility Header




Fig.3: NEMO Mobile Network Prefix Option




Fig.4: NEMO Mobile Network Prefix Length Option

3 BENEFITS OF MULTIHOMING

Currently, the term *multihoming* has no formal definition in the IETF community. It is used to indicate a situation where several routes are provided for network nodes to reach a particular correspondent. In the NEMO terminology, a mobile network is considered multihomed when either the mobile network is simultaneously connected to the Internet via more than one *mobile router*, or when a *mobile router* has more than one egress interface. Multihoming yields the following benefits:

- **Fault-Tolerance/Redundancy** is defined as the behavior in which the functions of a network are assumed by secondary system components when the primary component becomes unavailable (e.g. failure). As long as at least one connection to the Internet is maintained, the connectivity for all MNNs is guaranteed. This behavior can be split into two sub-classes:
 - Without transparency: The lost of one connection breaks ongoing transport sessions that use it but new transport sessions can be established.
 - With transparency: The lost of one connection is transparent to layers above the network layer, *i.e.* the lost of one connection does not disrupt ongoing transport sessions.
- **Load-Sharing** is defined as the spread of network traffic load among several routes. This is achieved when traffic load is distributed simultaneously among different connections between the mobile network and the Internet. Here, we do not indulge into forms of load balancing such as round-robin, least load first [23], *etc...* Load-sharing requires to set up multiple tunnels that will be used *simultaneously*. Load-sharing can be established either *statically* or *dynamically*.
- **Policy-Routing** is defined as the ability for the user or the application to choose between interfaces for matter of cost, efficiency, real-time politics, etc. This could be established either *statically* or *dynamically*, and initiated either by the MR, the HA, or a MNN.

4 NEMO BASIC SUPPORT ANALYSIS

The IETF NEMO working group is requiring *NEMO Basic Support* not to prevent multihoming [10, Section 5]. The purpose of this section is to evaluate NEMO Basic Support's ability to fulfill this requirement. We use the following taxonomy, originally proposed in the IETF NEMO WG in [20, Section 2], to classify all the potential multihoming configurations. Three discriminant factors, arbitrary named x, y, z, are used to differentiate these configurations:

- x: Indicates the number of MRs (presumably with multiple egress interfaces).
- y: Indicates the number of HAs associated with the mobile network.
- z: Indicates the number of MNPs announced to MNNs.

Eight cases have been identified using the tuple (x,y,z). A value of 1 implies there is single instance whereas a value of n indicates several instances. For each case, we detail the issues and the mechanisms required by NEMO Basic Support to allow redundancy, load sharing, and policy routing.

4.1 Case 1,1,1 Single mobile router, single home agent, single mobile network prefix




Fig.5: Case 1,1,1

This case is one of the most typical multihomed configuration. The mobile network is multihomed when the MR has multiple egress interfaces. Each egress interface of the MR is associated with a CoA. A bi-directional tunnel must thus be established between each CoA and the HA.

Redundancy: To maintain connectivity once a bi-directional tunnel is broken or disrupted, the MR only needs to transmit the traffic via another bi-directional tunnel set up on the same MR. This can be performed transparently to the MNNs because packet are always transmitted to/from the same MR's ingress interface, – *i.e.* independently of MR's links connectivity status. Such a behavior to provide redundancy for mobile nodes has been studied in [16]. From a NEMO Basic Support specification perspective, there doesn't seem to be any issue, but efficient link failure detection mechanisms provided by MR's hardware and software would be needed.

Load-Sharing: To provide for load sharing, NEMO Basic Support must at least:

1. Allow the simultaneous use of multiple bi-directional tunnels between one MR and one HA.
2. Allow to bind multiple CoAs to the same MNP.
3. Provide a mechanism to identify which CoA a Prefix-BU is meant to update.

In the NEMO Basic Support specification each egress interface can have its own HoA; thus each CoA can be bound to a specific HoA. However, the specification fails when the MR

has only one HoA. To illustrate the problem, the following entries would be expected to be recorded in HA's binding cache when the MR has two care-of addresses CoA-1 and CoA-2:

MNP/Prefix Length	Care-of Address
MNP-1/Length-1	CoA-1
MNP-1/Length-1	CoA-2

Following the current NEMO Basic Support specification, a Prefix-BU would only contain the following information:

MNP/Prefix Length	Care-of Address
MNP-1/Length-1	CoA-New

As we see, a way to differentiate between the CoAs the Prefix-BU is ought to update is needed. One solution is to use an extra identifier such as proposed in [25] [26].

To provide for *dynamic load-sharing*, additional information must be sent to the HA for each CoA. An appropriate place would be a load-sharing sub-option in the Prefix-BU to dynamically fix a *priority* level for each CoA. The HA would then be able to choose which CoA to send the traffic to according to this priority. However, such option is currently missing in NEMO Basic Support.

Policy-Routing For a dynamic management of the traffic based on preferences, policies associated with each CoA must be indicated to the HA for inbound traffic and to the MR for outbound traffic. NEMO Basic Support does not currently allow to carry such information. Several mechanisms may be contemplated: a policy field or a policy sub-option could be transported in the Prefix-BU between MR and HA by means of the NEMO Basic Support protocol, or a special policy exchange protocol could be designed. The *Flow Label* field in the IPv6 header [5, Section 6.] may be useful to identify which packet should be directed on which interface.

4.2 Case n,1,1: Multiple mobile routers, single home agent, single mobile network prefix




Fig.6: Case n,1,1

In this case, a bi-directional tunnel must be established between each MR and the HA. From a network mobility point of view, this case is similar to the previous one, thus the same requirements and recommendations apply. However, in the present case, each MR has only one route to the Internet through its own bi-directional tunnel. The distribution of outbound traffic is thus no longer operated by the MR between its multiple CoAs but by the MNNs between the MRs. Actually,

each MNN can decide through which MR it wants to send its packets using the router selection mechanism described in [19, Section 6.3.6]. So, no additional mechanisms are necessary to provide for redundancy with transparency. Moreover, an approach to manage the distribution of outbound traffic for Load-Sharing and Policy-Routing can be found in [8].

4.3 Case 1,n,1: Single mobile router, multiple home agents, single mobile network prefix




Fig.7: Case 1,n,1

In this case, the mobile network is reachable through several HAs, so the MR sends a Prefix-BU to each HA to establish a bi-directional tunnel on each of its multiple egress interfaces. We can define two sub-cases:

- HAs are in the same domain: Each HA advertises the same MNP via the routing protocol in the domain.
- HAs are in different domains: Due to complexity and lack of foreseen applications, this configuration is not advocated by the NEMO working group.

Each route from a HA to the MNP is announced in the fixed network by the corresponding HA. Inbound traffic is thus routed up to the tunnel entry by means of conventional routing protocols running in the fixed network, so redundancy shall easily be provided without additional support. In addition, a metric/cost can be recorded in the MNP advertisements. This could be configured statically at the HA or specified dynamically by the MR to each HA. However, in most cases, no routing protocol is running between the MR and its HAs, so link preferences cannot be propagated by the MR unless a preference option is specified in NEMO Basic Support.

4.4 Case n,n,1: Multiple mobile routers, multiple home agents, single mobile network prefix




Fig.8: Case n,n,1

In this case each MR independently maintains its own bidirectional tunnel. If a chain MR-tunnel-HA fails, outbound traffic can be moved to the other MR by default router selection [19, Section 6.3.6] whereas inbound traffic can be moved to the other HA by routing protocols. As illustrated earlier, both of these mechanisms are not specific to network mobility, thus no additional features are required in NEMO Basic Support.

4.5 Case 1,1,n: Single mobile router, single home agent, multiple mobile network prefixes




Fig.9: Case 1,1,n

In this case, distinct MNP are advertised in the mobile network. According to [24, Section 5.5] when an IPv6 node receives a router advertisement which contains several distinct network prefixes [19, Section 4.2], it auto-configures several addresses. Then each MNP performs a source address selection between addresses created with each mobile network prefix as defined in [7, Section 5]. Other multihoming issues and recommendations are the same than in case (1,1,1).

4.6 Other cases with multiple mobile network prefixes: Case n,1,n & 1,n,n & n,n,n




Fig.10: Case n,n,n

In these cases redundancy with transparency may not be ensured although there are multiple MRs or/and multiple HAs. Two reasons may account for disruption of service once a tunnels fails:

- Ingress filtering, for example reverse path filtering, may be made on source addresses by ISPs, HAs or even MRs. As a result, each tunnel may only carry the part of outbound traffic corresponding to one prefix, which creates a dichotomy between MNNs depending on their choice of source address.

- Source Address Selection: If a tunnel is broken, another one has to be used. In the event a MNP is not advertised anymore, MNNs must change their source address, which causes established sessions to break.

A solution proposed in [20] is to create a second binding on the ingress interface by sending a Prefix-BU through the other MR and then get several CoAs. By doing this, these cases translate respectively into cases n,1,1 or 1,n,1 or n,n,1. This requires more work. An analog mechanism can be found in [15]. Load-Sharing and Policy-Routing can be provided as detailed earlier.

4.7 Summary of our Analysis

In summary, the following mechanisms are missing in the NEMO Basic Support specification to allow all multihoming configurations:

- a mechanism to identify which CoA is updated by a Prefix-BU.
- a *preference* or *priority* sub-option in the Prefix-BU to provide for dynamic load-sharing.
- a *policy information* sub-option in the Prefix-BU for policy management.

However, nothing in the NEMO Basic Support specification seems to prevent the addition of the necessary missing mechanisms. Moreover, information necessary for load sharing and policy routing could be exchanged through a dynamic routing protocol running between the MR and the HA, or a special-purpose protocol.

5 CONCLUSION AND FUTURE WORK

Network mobility and multihoming are necessary to achieve the ubiquitous Internet we are waiting for. We thus studied NEMO Basic Support's ability to function under various multihoming configurations. All the contemplated scenarios can be classified into a taxonomy which makes the solution's ability easier to analyze. With this taxonomy, we were able to conclude that the specification of NEMO Basic Support does not prevent any of the multihoming configurations, but that some additional mechanism are requested if we want to provide redundancy, load-sharing, and policy-routing. Earlier studies for managing multiple interface for mobile hosts could indeed be helpful [25], [16], [3], [18], [26].

Based on the conclusions of our analysis detailed in [4] and other discussions on the NEMO mailing list, the NEMO WG has decided to add multihoming as a new working group item. The first task of the working group will likely be to determine what scenarios are useful to support from a deployment point of view, then to examine how the necessary mechanisms could be brought as extensions to NEMO Basic Support.

From our side, we are conducting further work on this subject within the Nautilus working group at WIDE [2]. The actual testing of these configurations remains to be done in

order to validate NEMO Basic Support. We are currently implementing NEMO Basic Support. Once this is done, we will validate the behavior of the protocol on our platform designed to test all the multihoming configurations, including nested multihomed mobile networks[17]. Based on our results, we will work on the necessary extensions so that all the expected benefits of multihoming can actually be enjoyed.

ACKNOWLEDGMENTS

We are grateful for comments on this paper provided by Ryuji Wakikawa (Keio University, Japan) and Nicolas Montavont (Université Louis Pasteur Strasbourg, France).

REFERENCES

- [1] NEMO web page and mailing list, June 2002. <http://www.ietf.org/html.charters/nemo-charter.html>.
- [2] NAUTILUS web page, April 2003. <http://www.nautilus6.org>.
- [3] F. Andre, J.-M. Bonnin, B. Deniau, K. Guillouard, N. Montavont, T. Noel, and L. Suci. Optimized Support of Multiple Wireless Interfaces within an IPv6 End-terminal. In *Smart Objects Conference (SOC)*, Grenoble, France, May 2003.
- [4] Julien Charbon and Chan-Wah Ng. Evaluating Multihoming Support in NEMO Basic Solution. *IETF*, July 2003. draft-charbon-nemo-multihoming-evaluation-00.txt.
- [5] Stephen E. Deering and Robert M. Hinden. Internet Protocol, Version 6 (IPv6) Specification. *IETF*, December 1998. RFC-2460.
- [6] Vijay Devarapalli, Ryuji Wakikawa, Alexandru Petrescu, and Pascal Thubert. NEMO Basic Support Protocol. Internet Draft draft-ietf-nemo-basic-support-00.txt, June 2003. Work in progress.
- [7] Richard Draves. Default Address Selection for Internet Protocol version 6 (IPv6). *IETF*, February 2003. RFC-3484.
- [8] Richard Draves and Robert M. Hinden. Default Router Preferences, More-Specific Routes, and Load Sharing. *IETF*, June 2002. draft-ietf-ipv6-router-selection-02.txt.
- [9] Thierry Ernst. *Network Mobility Support in IPv6*. PhD thesis, Université Joseph Fourier, October 2001. <http://www.inria.fr/rrrt/tu-0714.html>.
- [10] Thierry Ernst. Network Mobility Support Goals and Requirements. *IETF*, May 2003. draft-ietf-nemo-requirements-01.txt.
- [11] Thierry Ernst and Hong-Yon Lach. Network Mobility Support Terminology. *IETF*, May 2003. draft-ietf-nemo-terminology-00.txt.
- [12] Thierry Ernst, Koshiro Mitsuya, and Kei Uehara. Network Mobility from the InternetCAR Perspective. *JOIN: Journal on Interconnection Networks*, September 2003.
- [13] Thierry Ernst, Alexis Olivereau, Ludovic Bellier, Claude Castelluccia, and Hong-Yon Lach. Mobile Networks Support in Mobile IPv6 (Prefix Scope Binding Updates). Internet Draft draft-ernst-mobileip-v6-network-03.txt, March 2002. Work in progress.
- [14] Thierry Ernst and Keisuke Uehara. Connecting Automobiles to the Internet. In *ITST: 3rd International Workshop on ITS Telecommunications*, Seoul, South Korea, November 2002.
- [15] Jun-Ichiro Itojun Hagino and Hal Snyder. IPv6 Multihoming Support at Site Exit Routers. *IETF*, October 2001. RFC-3178.
- [16] N. Montavont, T. Noel, and M. Kassi-Lahlou. MIPv6 for Multiple Interfaces. Internet Draft draft-montavont-mobileip-mmi.txt, July 2002.
- [17] Nicolas Montavont, Thierry Ernst, and Thomas Noel. Multihoming in Nested Mobile Networks. In *International Symposium on Applications and the Internet (SAINT)- IPv6: Technology and Deployment Workshop*, Tokyo, Japan, January 2004.
- [18] Nicolas Montavont and Thomas Noel. A New Middleware for IPv6 Mobile Computing. In *International Conference on Networks (ICON)*, Sydney, Australia, September 2003.
- [19] Thomas Narten, Erik Nordmark, and William Allen Simpson. Neighbor Discovery for IP Version 6 (IPv6). *IETF*, December 1998. RFC-2461.
- [20] Chan-Wah Ng and Julien Charbon. Multi-Homing Issues in Bi-directional Tunneling. *IETF*, May 2003. draft-ng-nemo-multihoming-issues-01.txt.
- [21] Charles Perkins and David B. Johnson. Mobility Support in IPv6. In *Proc. of the Second*, pages 27–37, Rye, New-York, USA, November 1996. IBM Corporation and Carnegie Mellon University.
- [22] Pekka Savola. Examining Site Multihoming in Finnish Networks. Master’s thesis, Helsinki University of Technology, April 2003.
- [23] Pyda Srisuresh and Der-Hwa Gan. Load Sharing using IP Network Address Translation (LSNAT). *IETF*, August 1998. RFC-2391.
- [24] Susan Thomson and Thomas Narten. IPv6 Stateless Address Autoconfiguration. *IETF*, December 1998. RFC-2462.
- [25] Ryuji Wakikawa, Susumu Koshiba, Keisuke Uehara, and Jun Murai. Multiple Network Interfaces Support by Policy-Based Routing on Mobile IPv6. In *International Conference on Wireless Networks (ICWN)*, Las Vegas, USA, July 2002.
- [26] Ryuji Wakikawa, Keisuke Uehara, and Thierry Ernst. Multiple Care-of Address Registration on Mobile IPv6. Internet Draft draft-wakikawa-mobileip-multiplecoa-00.txt, February 2003.