

Network Working Group
Request for Comments: 4886
Category: Informational

T. Ernst
INRIA
July 2007

Network Mobility Support Goals and Requirements

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.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

Network mobility arises when a router connecting a network to the Internet dynamically changes its point of attachment to the Internet thereby causing the reachability of the said network to be changed in relation to the fixed Internet topology. Such a type of network is referred to as a mobile network. With appropriate mechanisms, sessions established between nodes in the mobile network and the global Internet can be maintained after the mobile router changes its point of attachment. This document outlines the goals expected from network mobility support and defines the requirements that must be met by the NEMO Basic Support solution.

Table of Contents

1. Introduction	2
2. NEMO Working Group Objectives and Methodology	3
3. NEMO Support Design Goals	5
3.1. Migration Transparency	5
3.2. Performance Transparency and Seamless Mobility	5
3.3. Network Mobility Support Transparency	5
3.4. Operational Transparency	5
3.5. Arbitrary Configurations	5
3.6. Local Mobility and Global Mobility	6
3.7. Scalability	7
3.8. Backward Compatibility	7
3.9. Secure Signaling	7
3.10. Location Privacy	8
3.11. IPv4 and NAT Traversal	8
3.12. Minimal Impact on Internet Routing	8
4. NEMO Basic Support One-Liner Requirements	8
5. Security Considerations	10
6. Acknowledgments	11
7. References	11
7.1. Normative References	11
7.2. Informative References	11

1. Introduction

Network mobility support (see [1] for the related terminology) is concerned with managing the mobility of an entire network, viewed as a single unit that changes its point of attachment to the Internet and thus its reachability in the Internet topology. Such a network is referred to as a mobile network and includes one or more mobile routers (MRs), which connect it to the global Internet. Nodes behind the MR(s) (MNNs) are both fixed (LFNs) and mobile (VMNs or LMNs). In most cases, the internal structure of the mobile network will be relatively stable (no dynamic change of the topology), but this is not always true.

Cases of mobile networks include, for instance:

- o Networks attached to people (Personal Area Networks or PANs): a cell phone with one cellular interface and one Bluetooth interface together with a Bluetooth-enabled PDA constitute a very simple instance of a mobile network. The cell phone is the mobile router while the PDA is used for web browsing or runs a personal web server.

- o Networks of sensors and computers deployed in vehicles: vehicles are increasingly equipped with a number of processing units for safety and ease of driving reasons, as advocated by ITS (Intelligent Transportation Systems) applications ([4]).
- o Access networks deployed in public transportation (buses, trains, taxis, aircrafts): they provide Internet access to IP devices carried by passengers (laptop, camera, mobile phone); host mobility within network mobility or PANs; network mobility within network mobility, i.e., nested mobility (see [1] for the definition of nested mobility).
- o Ad-hoc networks connected to the Internet via an MR: for instance, students in a train who need to both set up an ad-hoc network among themselves and get Internet connectivity through the MR connecting the train to the Internet.

Mobility of networks does not cause MNNs to change their own physical point of attachment; however, they do change their topological location with respect to the global Internet. If network mobility is not explicitly supported by some mechanisms, the mobility of the MR results in MNNs losing Internet access and breaking ongoing sessions between arbitrary correspondent nodes (CNs) in the global Internet and those MNNs located within the mobile network. In addition, the communication path between MNNs and correspondent nodes becomes sub-optimal, and multiple levels of mobility will cause extremely sub-optimal routing.

Mobility-related terms used in this document are defined in [2], whereas terms specifically pertaining to network mobility are defined in [1]. This document is structured as follows: in Section 2, we define the rough objectives and methodology of the NEMO working group to handle network mobility issues and we emphasize the stepwise approach the working group has decided to follow. A number of desirable design goals are listed in Section 3. Those design goals then serve as guidelines to define the requirements listed in Section 4 for basic network mobility support [3].

2. NEMO Working Group Objectives and Methodology

The mechanisms required for handling network mobility issues were lacking within the IETF standards when the NEMO working group (WG) was set up at the IETF in 2002. At that time, work conducted on mobility support (particularly in the Mobile IP working group) was to provide continuous Internet connectivity and optimal routing to mobile hosts only (host mobility support). Such mechanisms specified

in Mobile IPv6 [5] are unable to support network mobility. The NEMO working group has therefore been set up to deal with issues specific to network mobility.

The primary objective of the NEMO work is to specify a solution that allows mobile network nodes (MNNs) to remain connected to the Internet and continuously reachable while the mobile router serving the mobile network changes its point of attachment. The secondary goal of the work is to investigate the effects of network mobility on various aspects of Internet communication such as routing protocol changes, implications of real-time traffic and fast handovers, and optimizations. This should support the primary goal of reachability for mobile network nodes. Security is an important consideration too, and efforts should be made to use existing security solutions if they are appropriate. Although a well-designed solution may include security inherent in other protocols, mobile networks also introduce new challenges.

To complete these tasks, the NEMO working group has decided to take a stepwise approach. The steps in this approach include standardizing a basic solution to preserve session continuity (NEMO Basic Support, see [3]) and studying the possible approaches and issues with providing more optimal routing with potentially nested mobile networks (NEMO Extended Support, see [6] and [7] for a discussion on routing optimization issues and [8] for a discussion on multihoming issues). However, the working group is not chartered to actually standardize a solution for Extended Support at this point in time. If deemed necessary, the working group will be rechartered based on the conclusions of the discussions.

For NEMO Basic Support, the working group assumes that none of the nodes behind the MR is aware of the network's mobility; thus, the network's movement needs to be completely transparent to the nodes inside the mobile network. This assumption accommodates nodes inside the network that are not generally aware of mobility.

The efforts of the Mobile IP working group have resulted in the Mobile IPv4 and Mobile IPv6 protocols, which have already solved the issue of host mobility support. Since challenges to enabling mobile networks are vastly reduced by this work, basic network mobility support has adopted the methods for host mobility support used in Mobile IP and has extended them in the simplest way possible to achieve its goals. The Basic Support solution, now defined in [3] following the requirements stated in Section 4 of the present document, is for each MR to have a Home Agent (HA), and use bi-directional tunneling between the MR and HA to preserve session continuity while the MR moves. The MR acquires a Care-of Address (CoA) at its attachment point much like what is done for mobile hosts

(MHs), using Mobile IP. This approach allows nested mobile networks, since each MR will appear to its attachment point as a single node.

3. NEMO Support Design Goals

This section details the fundamental design goals the solutions will intend to achieve. Those design goals serve to define the issues and to impose a list of requirements for forthcoming solutions. Actual requirements for NEMO Basic Support are in Section 4; NEMO Extended Support is not yet considered at the time of this writing.

3.1. Migration Transparency

Permanent connectivity to the Internet has to be provided to all MNs, since continuous sessions are expected to be maintained as the mobile router changes its point of attachment. For maintaining those sessions, MNs are expected to be reachable via their permanent IP addresses.

3.2. Performance Transparency and Seamless Mobility

NEMO support is expected to be provided with limited signaling overhead and to minimize the impact of handovers on applications, in terms of packet loss or delay. However, although variable delays of transmission and losses between MNs and their respective CNs could be perceived as the network is displaced, it would not be considered a lack of performance transparency.

3.3. Network Mobility Support Transparency

MNs behind the MR(s) do not change their own physical point of attachment as a result of the mobile network's displacement in the Internet topology. Consequently, NEMO support is expected to be performed only by the MR(s). Specific support functions on any node other than the MR(s) would better be avoided.

3.4. Operational Transparency

NEMO support is to be implemented at the level of IP layer. It is expected to be transparent to upper layers so that any upper-layer protocol can run unchanged on top of an IP layer extended with NEMO support.

3.5. Arbitrary Configurations

The formation of a mobile network can occur in various levels of complexity. In the simplest case, a mobile network contains just a mobile router and a host. In the most complicated case, a mobile

network is multihomed and is itself a multi-level aggregation of mobile networks with collectively thousands of mobile routers and hosts. While the list of potential configurations of mobile networks cannot be limited, at least the following ones are desirable:

- o Mobile networks of any size, ranging from a sole subnet with a few IP devices to a collection of subnets with a large number of IP devices.
- o Nodes that change their point of attachment within the mobile network.
- o Foreign mobile nodes that attach to the mobile network.
- o Multihomed mobile network: either when a single MR has multiple attachments to the internet, or when the mobile network is attached to the Internet by means of multiple MRs (see definition in [1] and the analysis in [8]).
- o Nested mobile networks (mobile networks attaching to other mobile networks (see definition in [1])). Although the complexity requirements of these nested networks are not clear, it is desirable to support arbitrary levels of recursive networks. The solution should only impose restrictions on nesting (e.g., path MTU) when this is impractical and protocol concerns preclude such support.
- o Distinct mobility frequencies (see mobility factor in [2]).
- o Distinct access media.

In order to keep complexity minimal, transit networks are excluded from this list. A transit network is one in which data would be forwarded between two endpoints outside of the network, so that the network itself simply serves as a transitional conduit for packet forwarding. A stub network (leaf network), on the other hand, does not serve as a data forwarding path. Data on a stub network is either sent by or addressed to a node located within that network.

3.6. Local Mobility and Global Mobility

Mobile networks and mobile nodes owned by different administrative entities are expected to be displaced within a domain boundary or between domain boundaries. Multihoming, vertical and horizontal handoffs, and access control mechanisms are desirable to achieve this goal. Such mobility is not expected to be limited for any consideration other than administrative and security policies.

3.7. Scalability

NEMO support signaling and processing is expected to scale to a potentially large number of mobile networks irrespective of their configuration, mobility frequency, size, and number of CNs.

3.8. Backward Compatibility

NEMO support will have to co-exist with established IPv6 standards and not interfere with them. Standards defined in other IETF working groups have to be reused as much as possible and extended only if deemed necessary. For instance, the following mechanisms defined by other working groups are expected to function without modification:

- o Address allocation and configuration mechanisms.
- o Host mobility support: mobile nodes and correspondent nodes, either located within or outside the mobile network, are expected to continue operating protocols defined by the Mobile IP working group. This includes mechanisms for host mobility support (Mobile IPv6) and seamless mobility (FMIPv6).
- o Multicast support intended for MNNs is expected to be maintained while the mobile router changes its point of attachment.
- o Access control protocols and mechanisms used by visiting mobile hosts and routers to be authenticated and authorized, gaining access to the Internet via the mobile network infrastructure (MRs).
- o Security protocols and mechanisms.
- o Mechanisms performed by routers deployed in both visited networks and mobile networks (routing protocols, Neighbor Discovery, ICMP, Router Renumbering).

3.9. Secure Signaling

NEMO support will have to comply with the usual IETF security policies and recommendations and is expected to have its specific security issues fully addressed. In practice, all NEMO support control messages transmitted in the network will have to be protected with an acceptable level of security to prevent intruders from usurping identities and forge data. Specifically, the following issues have to be considered:

- o Authentication of the sender to prevent identity usurpation.

- o Authorization, to make sure the sender is granted permission to perform the operation as indicated in the control message.
- o Confidentiality of the data contained in the control message.

3.10. Location Privacy

Location privacy means hiding the actual location of MNN to third parties other than the HA are desired. It is not clear to which extend this has to be enforced, since it is always possible to determine the topological location by analyzing IPv6 headers. It would thus require some kind of encryption of the IPv6 header to prevent third parties from monitoring IPv6 addresses between the MR and the HA. On the other hand, it is at the very least desirable to provide a means for MNNs to hide their real topological location to their CNS.

3.11. IPv4 and NAT Traversal

IPv4 clouds and NAT are likely to co-exist with IPv6 for a long time, so it is desirable to ensure that mechanisms developed for NEMO will be able to traverse such clouds.

3.12. Minimal Impact on Internet Routing

Any NEMO solution needs have minimal negative effect on the global Internet routing system. The solution must therefore limit both the amount of information that must be injected into Internet routing, as well as the dynamic changes in the information that is injected into the global routing system.

As one example of why this is necessary, consider the approach of advertising each mobile network's connectivity into BGP and, for every movement, withdrawing old routes and injecting new routes. If there were tens of thousands of mobile networks each advertising and withdrawing routes, for example, at the speed that an airplane can move from one ground station to another, the potential effect on BGP could be very unfortunate. In this example, the total amount of routing information advertised into BGP may be acceptable, but the dynamic instability of the information (i.e., the number of changes over time) would be unacceptable.

4. NEMO Basic Support One-Liner Requirements

For basic network mobility support, the NEMO WG is to specify a unified and unique "Network Mobility (NEMO) Basic Support" solution, hereafter referred to as "the solution". This solution is to allow all nodes in the mobile network to be reachable via permanent IP

addresses, as well as maintain ongoing sessions as the MR changes its point of attachment to the Internet topology. This is to be done by maintaining a bi-directional tunnel between an MR and its Home Agent.

The NEMO Working Group, after some investigation of alternatives, has decided to reuse and extend the existing Mobile IPv6 [5] mechanisms for tunnel management.

The list of requirements below has been imposed on the NEMO Basic Support solution. The requirements have mostly been met by the resulting specification, which can now be found in [3]. Associated deployment issues are discussed in [9].

R01: The solution must be implemented at the IP layer level.

R02: The solution must set up a bi-directional tunnel between a mobile router and its Home Agent (MRHA tunnel).

R03: All traffic exchanged between an MNN and a CN in the global Internet must transit through the bi-directional MRHA tunnel.

R04: MNNs must be reachable at a permanent IP address and name.

R05: The solution must maintain continuous sessions (both unicast and multicast) between MNNs and arbitrary CNs after IP handover of (one of) the MRs.

R06: The solution must not require modifications to any node other than MRs and HAs.

R07: The solution must support fixed nodes, mobile hosts, and mobile routers in the mobile network.

R08: The solution must allow MIPv6-enabled MNNs to use a mobile network link as either a home link or a foreign link.

R09: The solution must ensure backward compatibility with other standards defined by the IETF. In particular, this includes the following:

R09.1: The solution must not prevent the proper operation of Mobile IPv6 (i.e., the solution must allow MIPv6-enabled MNNs to operate either the CN, HA, or MN operations defined in [5]).

- R10: The solution must be agnostic to the internal configuration. This means the solution will behave the same way if NEMO is nested, comprises one or several subnets, or comprises MNs that are LFNs, VMNs, LFNs or a mixture of them.
- R11: The solution must support at least 2 levels of nested mobile networks, while, in principle, arbitrary levels of recursive mobile networks should be supported.
- R12: The solution must function for multihomed MRs and multihomed mobile networks as defined in [1].
- R13: NEMO support signaling over the bi-directional must be minimized.
- R14: Signaling messages between the HA and the MR must be secured:
- R14.1: The receiver must be able to authenticate the sender.
 - R14.2: The function performed by the sender must be authorized for the content carried.
 - R14.3: Anti-replay must be provided.
 - R14.4: The signaling messages may be encrypted.
- R15: The solution must ensure transparent continuation of routing and management operations over the bi-directional tunnel (this includes, e.g., unicast and multicast routing protocols, router renumbering, Dynamic Host Configuration Protocol (DHCPv6)).
- R16: When one egress interface fails, the solution may preserve sessions established through another egress interface.
- R17: The solution should have a minimal impact on the global Internet routing system.

5. Security Considerations

Security considerations of the NEMO Basic Support solution are addressed in [RFC3963].

Section 3.9 of this document discusses the security goals for all forms of existing and forthcoming NEMO solutions.

6. Acknowledgments

The material presented in this document takes most of its text from discussions and previous documents submitted to the NEMO working group. This includes initial contributions from Motorola, INRIA, Ericsson, and Nokia. We are particularly grateful to Hesham Soliman (Ericsson) and the IETF Area Directors (ADs) at the time (Erik Nordmark and Thomas Narten), who greatly helped to set up the NEMO working group. We are also grateful to all the following people whose comments greatly contributed to the present document: T.J. Kniveton (Nokia), Alexandru Petrescu (Motorola), Christophe Janneteau (Motorola), Pascal Thubert (Cisco), Hong-Yon Lach (Motorola), Mattias Petterson (Ericsson), and all the others who have expressed their opinions on the NEMO mailing lists (formerly known as MONET). Thierry Ernst wishes to personally acknowledge INRIA Rhone-Alpes and Motorola Labs Paris for their support and direction in bringing up this topic to the IETF in 2001--particularly Claude Castelluccia (INRIA) and Hong-Yon Lach (Motorola)--and his past employer, Keio University, Japan, which supported most of the costs associated with the IETF during the timelife of previous versions of this document.

7. References

7.1. Normative References

- [1] Ernst, T. and H. Lach, "Network Mobility Support Terminology", RFC 4885, July 2007.
- [2] Manner, J. and M. Kojo, "Mobility Related Terminology", RFC 3753, June 2004.
- [3] Devarapalli, V., Wakikawa, R., Petrescu, A., and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", RFC 3963, January 2005.

7.2. Informative References

- [4] "CALM - Medium and Long Range, High Speed, Air Interfaces parameters and protocols for broadcast, point to point, vehicle to vehicle, and vehicle to point communication in the ITS sector - Networking Protocol - Complementary Element", ISO Draft ISO/WD 21210, February 2005.
- [5] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [6] Ng, C., Thubert, P., Watari, M., and F. Zhao, "Network Mobility Route Optimization Problem Statement", RFC 4888, July 2007.

- [7] Ng, C., Zhao, F., Watari, M., and P. Thubert, "Network Mobility Route Optimization Solution Space Analysis", RFC 4889, July 2007.
- [8] Ng, C., Ernst, T., Paik, E., and M. Bagnulo, "Analysis of Multihoming in Network Mobility Support", Work in Progress), February 2007.
- [9] Thubert, P., Wakikawa, R., and V. Devarapalli, "Network Mobility (NEMO) Home Network Models", RFC 4887, July 2007.

Author's Address

Thierry Ernst
INRIA
INRIA Rocquencourt
Domaine de Voluceau B.P. 105
78 153 Le Chesnay Cedex
France

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

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.

Acknowledgement

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

