NEMO Working Group M. Tsukada
Internet-Draft R. Kuntz
Expires: April 20, 2006 T. Ernst
Keio University / WIDE
October 17, 2005

Analysis of Multiple Mobile Routers Cooperation
draft-tsukada-nemo-mr-cooperation-analysis-00.txt

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with Section 6 of BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt.

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

This Internet-Draft will expire on April 20, 2006.

Copyright Notice

Copyright (C) The Internet Society (2005).

Abstract

This document is an analysis of multiple Mobile Routers Cooperation in the context of network mobility support (NEMO) in IPv6. Our objective is to identify when cooperation between MRs is needed and what information must be exchanged.

Table of Contents

1. Introduction 3

2. Objectives, Motivations and Assumptions 3
2.1 Objectives 3
2.2 Sample Scenarios 4
2.2.1 Mobile Network in a Personal Area Network 4
2.2.2 Mobile Network in a Vehicle 5
2.3 Benefits 5
2.4 Assumptions 6

3. Points Of Failure 7
3.1 MR Failure (B) 8
3.2 HA Failure (D) 8
3.3 NEMO Tunnel Failure (C) 8
3.4 NEMO-Link Failure (A) 9

4. Requirements Analysis 9
4.1 NEMO Tunnel Discovery 10
4.2 Ingress Filtering Discovery 11
4.3 MR State Discovery 12
4.4 Link Characteristics Discovery 13
4.5 MNP Relay 13
4.6 Split NEMO Discovery 13
4.7 Policy Discovery 13

5. Approaches 14
5.1 Mobility Approach 14
5.1.1 Binding Information Sharing 14
5.1.2 Link Characteristics Sharing 15
5.1.3 Policy Sharing 16
5.2 Standard approach 16

6. Security Consideration 16

7. Conclusion 16

8. Acknowledgement 16

9. References 17

Authors' Addresses 17

Intellectual Property and Copyright Statements 19

1. Introduction

Network mobility support is necessary for groups of computers moving together and requiring access to the Internet, such as networks of sensors or access networks deployed in vehicles. NEMO Basic Support (or NEMO in short [2]) has been specified to address this need.

In some situations, the mobile network may be multihomed, that is either when a mobile router is multihomed or when several mobile routers are being used to attach the same mobile network to the Internet [5]. This brings a number of benefits [6] including for instance the possibility to face the lack of coverage of a particular technology, to increase the Internet connectivity and to choose the best path in terms of delay, bandwidth or price.

In theory, the simultaneous use of multiple Mobile Routers (MRs) improves the overall Internet connectivity offered to the Mobile Network Nodes (MNN) located in the mobile network, because multiple paths to the Internet are available. However, in practice, it cannot be realized using NEMO basic support only. One of the reasons is that an MR in a mobile network does not know when another Internet connectivity (provided by another MR) is available.

Some of the issues are outlined in [5] and cooperation between MRs is advocated, but we need to further detail the problem before we can bring a solution for such a cooperation between MRs.

At first we describe in section Section 2 the motivations and scenarios for using multiple MRs. Then, in Section 3 we detail the problems or limitations caused by the use of multiple MRs to connect a NEMO to the Internet. From this problem description, we are able to list a number of requirements in Section 4 and finally we discuss in Section 5 potential approaches for MR cooperation.

2. Objectives, Motivations and Assumptions

2.1 Objectives

The objectives of this document are:

- To analyze further the cases where multiple MRs are available in the NEMO (cases (n,*,*) as categorized in [5]).
- To pave the way for a comprehensive solution to the issues applying to such a multihomed NEMO configuration.
- To clarify the MR cooperation issues and to list the required mechanisms for MR cooperation.

- To discuss possible approaches.

2.2 Sample Scenarios

2.2.1 Mobile Network in a Personal Area Network

In a Personal Area Network (PAN), both a mobile phone and a PDA are acting as MRs and connecting the mobile network to the Internet (Figure 1). The MRs have different access technologies. A GPRS interface is provided by the mobile phone, and an IEEE802.11g interface is shipped with the PDA. Some MNNs such as a portable game player and a small PC are connected to the PAN. These MNNs may communicate with Correspondent Nodes (CN) with different needs in terms of cost, efficiency, bandwidth requirement, delay, etc. For example, a network game played on the portable game player may require less delay and higher bandwidth than browsing on the PC. In this case, the traffic of the network game could use IEEE802.11g access available on the PDA, whereas the web access could use both GPRS and IEEE802.11g. The use of multiple MRs simultaneously allows to increase the bandwidth and to improve the access redundancy for the MNNs located in the mobile network. However, the current NEMO Basic Support does not provide any solution to share the connectivities among several MRs. We thus consider that there is a need for MRs cooperation.

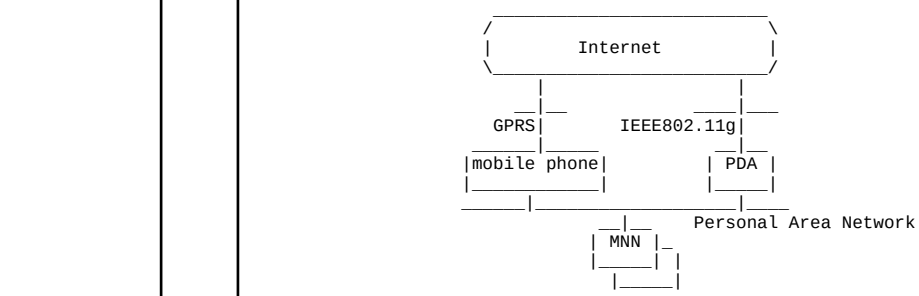

Figure 1: Mobile Network in a Personal Area Network

2.2.2 Mobile Network in a Vehicle

We are assuming network of sensors and computers deployed in vehicles as shown in Figure 2. There are several MNNs in a vehicle and the vehicle network is connected to the Internet via two in-vehicle MRs. One is a powerful MR that provides high-bandwidth to the MNN. As it may consume some energy, it is only switched on when the car is started. When the car is turned off, this MR is switched off. The other MR is a low-power, low-cost and low-bandwidth MR that is used as a backup while the vehicle is stopped (ie when the engine is off) or when the main MR sustains some failures. The use of multiple MRs allows to access anytime the sensible information of the car. Some cooperation between those MRs would allow to ensure that a path is always available from and to the MNNs.

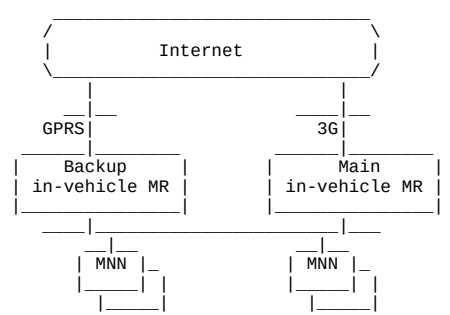

Figure 2: Mobile Network in a Vehicle

2.3 Benefits

The benefits of Mobile Routers cooperation in the (n,*,*) cases are the same as the multihoming benefits detailed in [6]:

1. Permanent and Ubiquitous Access / Redundancy / Fault-Recovery: When a tunnel fails on an MR, MNNs can not communicate with their correspondents, and conversely. If MRs cooperate with each other, communication can be recovered via the other MR relaying the failed one. There are thus more chances to maintain a permanent connectivity to the Internet.

2. Load Sharing / Load Balancing / Preference Settings / Increased Bandwidth / Bi-casting: When several tunnels are available simultaneously, traffic can be spread among several MRs, either by sharing the communication flows between MRs, or bicasting (n-casting) packets. In addition, and provided enough information about the available paths and their characteristics (link efficiency, cost, etc), packets could be routed according to preferences set up by the users or the applications.

2.4 Assumptions

Following the taxonomy proposed in [5], we will concentrate the analysis taking into account all the topologies where multiple MRs are available in the same NEMO. This includes all the (n,*,*) cases, namely:

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

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

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

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

We assume that tunnels are set up using NEMO Basic Support and some extensions to support multiple binding registrations, as currently investigated by the Monami6 WG. Three configurations can be divided when multiple MRs are connecting the same NEMO to the Internet, as shown in Figure 3. These MRs connecting the same NEMO are referred to as "neighbor MRs" in some parts of the present document.

In the (n,1,*) case, all (MR,HA) tunnels end up at the same HA. In the (n,n,*) case, either all (MR,HA) tunnels may end up at different HAs in a "pair tunnel" type or tunnels may be established between each MR and several HAs in "mesh like" type.

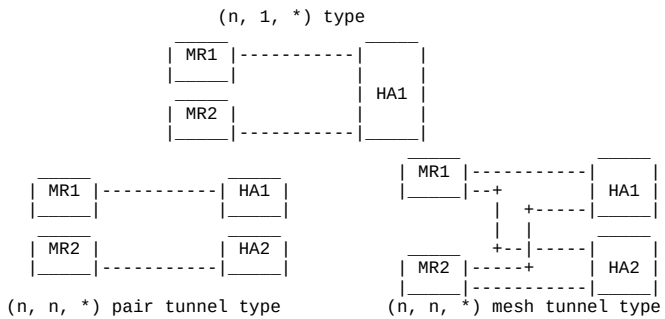

Figure 3: Tunnel Topology Type

Nested NEMOs are out of scope of the present study, and we don't make any assumption about the specific NEMO topology. MRs may be connected to one another on a single NEMO-link, but the NEMO could also be made of several NEMO-links, with MRs on independent NEMO-links.

3. Points Of Failure

Some of the issues are outlined in [5] and cooperation between MRs is recommended, but we need to further detail the problem before we can bring a solution for such a cooperation between MRs. We will make this study from a failure point of view, i.e. by investigating the potential points of failure and remedy to act upon these failures.

There are four points of failure that may affect the way multiple MRs are supported. These are illustrated on Figure 4: (A) on the NEMO-link, (B) on the MR, (C) on the (MR,HA) tunnel and (D) on the HA.

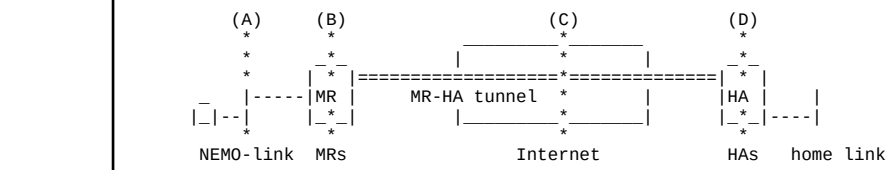

Figure 4: Failure Analysis

3.1 MR Failure (B)

The failure of the MR could happen when the MR is down or when it is isolated from the other MRs. (e.g. all interfaces are down). As a result, the failed MR becomes unreachable and thus unable to communicate with other MRs and HAs. All tunnels established by this MR would be disrupted (see point (B) on Figure 4). In addition, the failure may affect the MR which was advertising an MNP, which may cause some MNNs to renumber.

A mechanism is needed to detect the failure of the MR and to redirect the traffic over an alternative tunnel (possibly newly established) between another MR and the same or another HA. It may also be necessary to relay the MNP that was advertised by the failed MR.

3.2 HA Failure (D)

The failure of the HA could happen when the HA is down or when it is isolated from the Internet, for instance when the access network where the interface that connects it to the Internet fails or if the fault occurs at the edge of the network. As a result, the HA becomes unreachable (see point (D) on Figure 4).

As a result, the failed HA is unable to communicate with MRs and other HAs. The HA can not be used to register the Care-of Addresses, all tunnels established with this HA would be disrupted, and packets could not be redirected from the CNs to the mobile network anymore. In addition, the MNPs registered by the MRs on this HA become unreachable, and the MRs cannot return home.

A mechanism is needed to detect the failure of the HA and to redirect the traffic over an alternative tunnel (possibly newly established) between the same MR and another HA.

3.3 NEMO Tunnel Failure (C)

Events happening between the MR and its HA (see point (C) on Figure 4) could cause the tunnel to be disrupted: e.g. one egress interface of the MR may fail, the access router may fail, the access network may be disconnected from the Internet, or the home link may fail.

The MR can not send Binding Updates nor receive incoming packets from the HA through this specific tunnel. However the MR may still be able to communicate with its peer HAs via other paths and other

tunnels could established on another egress interface or via the other MR.

A mechanism is needed to detect the tunnel failure and to redirect all the traffic over an alternative tunnel (possibly newly established), between the same or another (HA, MR) pair.

3.4 NEMO-Link Failure (A)

Events happening between two MRs on the same NEMO may cause two MRs to be disconnected: e.g. the ingress interface of an MR may fail, or the link connecting the two MRs may fail (a cable or hub if Ethernet is used, of some radio interferences when a wireless technology is used).

Each MR may continue to send Binding Updates to its HA. No established tunnel may be affected but the failure may cause the NEMO to split. As a result, the incoming packets to the mobile network may not reach the MNN the packet is bound to. Also, Multiple MRs connected to this NEMO-Link will not be able to communicate via this link anymore.

A mechanism is therefore needed to detect the link failure and to act upon it.

4. Requirements Analysis

Path selection in multihomed NEMO is complex because there are several steps to determine a path. We thus classify the steps for path selection as follows (see Figure 5): (A) MNNs select the exit MRs as default routers, (B) the MRs select a path to the HAs, (C) the HAs select a path to the mobile network and (D) the CN select a path to the HAs.

In the face of the failures highlighted previously, the path may need to be redirected from an MR to another. If the failure happens on the HA or the NEMO Tunnel, this redirection may be done in cooperation between the former MR to the new MR. If the failure happens on the MR or the NEMO-link, this redirection may be done by the new MR alone.

The MNNs may send the packets to their default MR or they may have the ability to choose the MR, but in any case, the ultimate choice of the exit bi-directional tunnel falls on MRs hands. This choice could be made based on the source and destination address in the packet, or based on other parameters such as the availability of tunnels. For such selection to be made efficiently, MRs need to share information about the tunnels existing on the other MRs, the state of the MR, and

other characteristics.

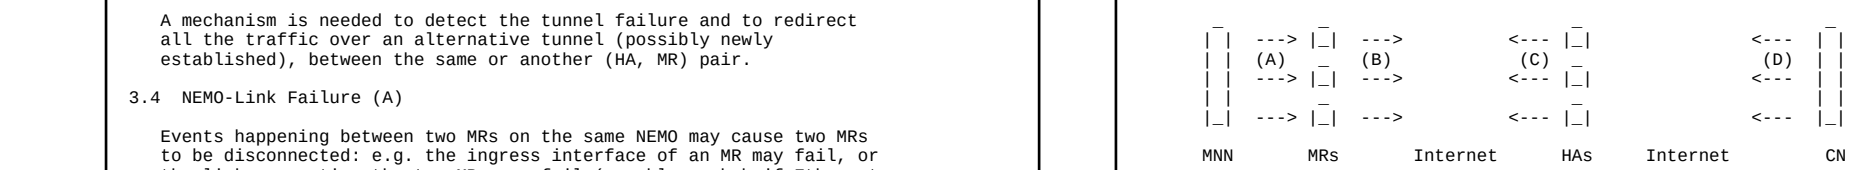

Figure 5: Path Selection

The exchange of information will help each MR to determine the network conditions and available Internet accesses. As the NEMO and tunnel infrastructures may change over time, it is mandatory that such information be regularly exchanged or that the exchange is triggered by an MR when a change is detected.

4.1 NEMO Tunnel Discovery

Each MR must be aware of the tunnels established at other MRs with their respective HAs in order to know which paths are available to route packets to the Internet. Both MRs and HAs should be aware of all NEMO tunnels between MRs and a HA(s) irrespectively of the administrative domain of the HA. A NEMO tunnel discovery mechanism is therefore required.

For example, in Figure 6 below, let’s see what tunnels are seen simply using NEMO Basic Support:

- In the (n,1,*) type, HA1 would be aware of tunnel (A) and (B), whereas MR1 would only be aware of tunnel (A) and MR2 of tunnel (B). MR1 is unaware of tunnel (B) and MR2 is unaware of tunnel (A).

- In the (n,n,*) pair tunnel type, MR1 and HA1 would be aware of tunnel (C) and MR2 and HA2 would be aware of tunnel (D). However, MR1 and HA1 are unaware of tunnel (D) and MR2 and HA2 are unaware of tunnel (C).

- In the (n,n,*) mesh tunnel type, MR1 would be aware of tunnels (E) and (G) whereas MR2 would be aware of tunnel (F) and (H). MR1 is unaware of tunnel (F) and (H) whereas MR2 is unaware of tunnels (E) and (G).

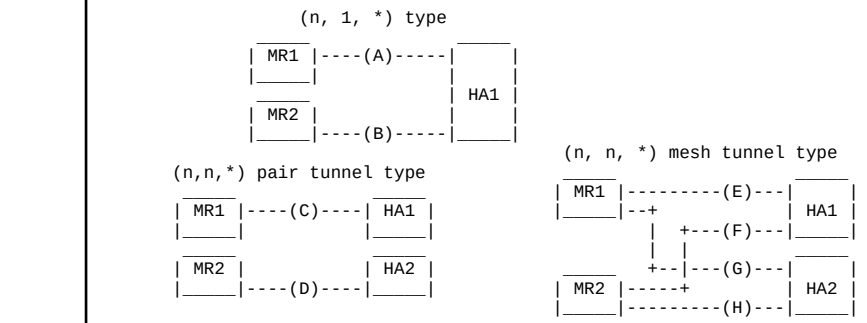

Figure 6: NEMO Tunnel Discovery

The tunnels which need to be discovered are summarized in the matrix below. Mark [X] stands for a node (MR or HA) that cannot be aware of the tunnel only using NEMO Basic Support.

(n,1,*) type			(n,n,*) pair tunnel type			(n,n,*) mesh tunnel type				
	(A)	(B)		(C)	(D)		(E)	(F)	(G)	(H)
MR1		X	MR1		X	MR1	X		X	
MR2	X		MR2	X		MR2		X		X
HA1			HA1		X	HA1	X	X		
			HA2	X		HA2			X	X

Figure 7: NEMO Tunnel Discovery Matrix

4.2 Ingress Filtering Discovery

Due to ingress filtering (see [5]), each MR has to know which address range can be authorized over a specific (MR,HA) tunnel. For doing so, the knowledge of the multihomed NEMO configuration an MR belongs to, i.e. (x,y,z) case, and to know the (MR,HA) tunnel topology type

as described in Section 2.4 are necessary.

For example, in Figure 6 above, let's see what existing tunnels are valid:

- o In (n,n,n) pair tunnels type, MR1-HA1 and MR2-HA2 tunnels are available. Sessions between MNN and CN may not be directed from a tunnel to the other without breaking sessions because both tunnels end up at a different HA.
- o In (n,n,n) mesh tunnels type, all possible tunnels, MR1-HA1, MR1-HA2, MR2-HA1 and MR2-HA2 tunnels are available. Sessions between MNN and CNs can be changed from MR1-HA1 tunnel to MR2-HA1 tunnel because both tunnels end up at a same HA, so ingress filtering doesn't occur. However, traffic may not be redirected to MR1-HA2 or MR2-HA2.

4.3 MR State Discovery

Each MR must be aware of its neighbor MRs' state. In case an MR is down, this would avoid one MR to attempt cooperation with a failed neighbor. This would also allow to trigger a relaying mechanism for the failed MR. In case an MR recovers from a failure, this would allow to return to the initial state of the NEMO, and to stop any relaying mechanisms that was set up before the MR failed.

By exchanging crucial information via the NEMO-link, each MR will be able to discover when a neighbor MR has failed or recovered, and to take some decisions to setup or stop a relaying mechanism.

The MR state discovery mechanism should not assume any specific configuration of the NEMO. It should work either for MRs on the same NEMO-link, or for MRs on distinct NEMO-links as shown on Figure 8. However, in the latter case, special care must be taken so that such discovery mechanism does not introduce loops or complexity.

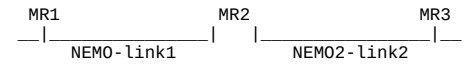

Figure 8: MRs Discovery Between Different NEMO-links

4.4 Link Characteristics Discovery

A MR usually connects to the access networks using wireless technologies that may have different characteristics (e.g. link quality, bandwidth, delay and cost). Sharing some of such link metrics among the MRs will help to perform a decision in the path selection mechanism. The following information can be interesting to select an access link on an MR:

- Available bandwidth / Used bandwidth
- Expected delay (RTT) to the HA
- SNR (Signal to Noise Ratio)
- Security mechanisms

The protocol used to exchange such metrics should be easily extensible as new wireless technologies may bring new interesting link metrics in the future.

4.5 MNP Relay

A recovery mechanisms is needed to continue advertising an MNP upon a failure happening on an MR, a HA or a NEMO-link which affects the advertisement of an MNP.

4.6 Split NEMO Discovery

When an MR notices that one of its neighbor MR is unreachable, it may be due to a failure or because the NEMO-link has split. A mechanism is thus necessary to be able to distinguish when a split occurs.

The mechanism should also solve the Prefix Ownership issue as described in [5].

4.7 Policy Discovery

Policies are used to divide flows among several available paths according to the flow type, the port number, the destination address etc. For the packets sent from the mobile network, multiple MRs need to work with the same policies in order to take the same and coherent decisions. Thus a mechanism to synchronize policies among the MRs is mandatory.

Policies should be exchanged using a standardized format among the MRs as they may use different OS or policy processing tools.

5. Approaches

In this section we investigate possible approaches to meet the requirements outlined in the previous section. One approach is to share mobility related information, and another to use standard protocols.

The way to exchange information among neighbor MR may be different if the MRs are located on the same link or in two different links (connected via a router or another MR). In the former case, sending the information to the all-router multicast address or to the targeted MR's link local address will allow the MRs to exchange information even if they are not on the same IP network. In the latter case, a new mechanism to exchange the information between the networks will be needed.

5.1 Mobility Approach

5.1.1 Binding Information Sharing

MRs that are in the same mobile network can share among them the state of their binding information as well as some other important items, namely:

- Care-of Address / Prefix Length

  By exchanging their Care-of Address(es) (with the prefix length), different MRs will able to detect that they are connecting to same access network. It is (usually) not interesting that several MRs connect to same access network regarding to redundancy, bandwidth or battery purposes. The MRs may able to turn off the interfaces that are connected to the same access network except for one interface.

- Home Address / Prefix Length

  By exchanging their Home Address(es) (with the prefix length), different MRs can know the home link of a specific MR, since the Home Address is built from the the MR's home link's prefix.

- Mobile Network Prefix / Prefix Length

  With the MNP and its prefix length, the MRs can know which MRs are advertising the same MNP to the NEMO. This could help to solve the Prefix Ownership issue (as described in [5]) when a mobile network where several MRs advertise the same MNP splits.

- o HA address

  With the HA address, MRs can know where a specific tunnel ends. Thanks to this information, when a communication has to be diverted from one tunnel to another, ingress filtering can be avoided by choosing efficiently the new tunnel.

- o Ingress interface's link-local address

  By exchanging their link local address, MRs on the same link can communicate directly via the NEMO-link, even if they are on two different IPv6 networks.

- o Ingress interface's global address

  By exchanging their global address, MRs that are on the same mobile network but not on the same link may be able to communicate via the NEMO-link.

- o Binding Lifetime

  With the binding lifetime, each MR knows when other's MRs binding information have expired and can thus easily remove deprecated information.

- o MR Identification (MR ID)

  MR should be distinguished by an unique identifier. This ID must not be duplicated in a same mobile network.

- o Tunnel Identification (tunnel ID)

  Tunnel Identification is necessary to be able to distinguish the multiple available path on a single MR. This ID must be unique on the same MR.

Those information can be maintained with the MR Identification and the Tunnel Identification as key of the information database.

5.1.2 Link Characteristics Sharing

The parameters exchanged should be bound to the tunnel ID and the MR ID.

[8] is a solution that falls into this approach. It attempts to propose a link metric message exchange solution which allows MNNs to select the MR with the best link quality.

5.1.3 Policy Sharing

Each policy exchanged should be bound to a tunnel ID and an MR ID.

The path for incoming packets to the mobile network may be selected by the HA. In this document we focus on MRs cooperation, however it is recommended to share the same solution about policy sharing for both the MR-MR and MR-HA cases.

5.2 Standard approach

The standard approach would use or extend existing protocols to fit the requirements, such as:

- o A routing protocol to announce and propagate the available path to the Internet inside the NEMO. Metrics such as described in Section 5.1.2 could also be exchanged in this way.

- o Neighbor Discovery to discover the neighbor MRs located in the same link.

- o A prefix delegation mechanism for the MRs to deal with the MNPs advertised in the NEMO.

6. Security Consideration

Security considerations are out of scope of the present document, since we do not discuss specific solutions.

7. Conclusion

In this document, we investigated the issues for NEMO configurations with multiple MRs. We have listed a number of requirements that should be met for MNNs and MRs to choose the most appropriate exit tunnel and to recover from a failure of the MR, the HA, the tunnel itself or the NEMO-link between MRs. We have started to investigate potential approaches and we outline one which is based on the exchange of mobility binding information, and a standard one using existing network routing protocols, router advertisement and prefix delegation mechanisms.

8. Acknowledgement

The authors would like to thank the Nautilus6 members for the initial discussions that has motivated the writing of this draft.

9. References

[1] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.

[2] Devarapalli, V., Wakikawa, R., Petrescu, A., and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", RFC 3963, January 2005.

[3] Ernst, T., "Network Mobility Support Goals and Requirements", draft-ietf-nemo-requirements-04 (work in progress), February 2005.

[4] Ernst, T. and H. Lach, "Network Mobility Support Terminology", draft-ietf-nemo-terminology-03 (work in progress), February 2005.

[5] Ng, C., Paik, E., Ernst, T., and M. Bagnulo, "Analysis of Multihoming in Network Mobility Support", draft-ietf-nemo-multihoming-issues-03 (work in progress), July 2005.

[6] Ernst, T., "Goals and Benefits of Multihoming", draft-ernst-generic-goals-and-benefits-01 (work in progress), February 2005.

[7] Wakikawa, R., Uehara, K., Ernst, T., and K. Nagami, "Multiple Care-of Addresses Registration", draft-wakikawa-mobileip-multiplecoa-04 (work in progress), June 2005.

[8] Morioka, H., "Mobile Router Cooperation Protocol", draft-morioka-nemo-mrcoop-00 (work in progress), July 2005.

Authors' Addresses

Manabu Tsukada
Keio University / WIDE
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: tu-ka@sfc.wide.ad.jp
URI: http://www.sfc.wide.ad.jp/~tu-ka/

Romain Kuntz
Keio University / WIDE
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: kuntz@sfc.wide.ad.jp
URI: http://www.sfc.wide.ad.jp/~kuntz/

Thierry Ernst
Keio University / WIDE
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/

Intellectual Property Statement

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.

Disclaimer of Validity

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 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.

Copyright Statement

Copyright (C) The Internet Society (2005). 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.

Acknowledgment

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