

MEXT Working Group  
Internet-Draft  
Intended status: Informational  
Expires: May 7, 2009

C. Bernardos  
M. Bagnulo  
UC3M  
November 3, 2008

**Analysis on how to address NEMO R0 for Aeronautics Mobile Networks  
draft-bernardos-mext-aero-nemo-ro-sol-analysis-01**

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 May 7, 2009.

Abstract

The Network Mobility Basic Support protocol enables networks to roam and attach to different access networks without disrupting the ongoing sessions that nodes of the mobile network may have. By extending the Mobile IPv6 support to Mobile Routers, nodes of the mobile network are not required to support any kind of mobility, since packets go through the Mobile Router-Home Agent (MRHA) bi-directional tunnel. Data packets belonging to communications of nodes of the mobile network have to traverse the Home Agent, and therefore resulting paths are likely to be suboptimal. Additionally, the solution adds packet overhead, due to the use of encapsulation between the Mobile Router and the Home Agent.

There are currently a set of well defined NEMO Route Optimization requirements for Operational Use in Aeronautics and Space Exploration, which potential solutions should meet. This document analyses how the problem of NEMO RO for Aeronautics Mobile Networks might be tackled, in a way as generic as possible, trying to identify those open questions and deployment considerations that need to be addressed.

The main goal of this document is to foster the discussion about aeronautics NEMO RO solution space within the MEXT WG.

#### Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC 2119](#) [1].

Table of Contents

- [1.](#) Introduction . . . . . [4](#)
- [2.](#) Solution Space analysis . . . . . [4](#)
- [3.](#) Design issues/questions/trade-offs . . . . . [8](#)
  - [3.1.](#) Where are the RO entities located? . . . . . [9](#)
  - [3.2.](#) Who administratively manages the RO entities? . . . . . [10](#)
  - [3.3.](#) Which kind of addresses are gonna be used and who own them? . . . . . [11](#)
  - [3.4.](#) How many RO entities are needed to globally perform NEMO RO? . . . . . [11](#)
  - [3.5.](#) What trust relationships are needed? . . . . . [12](#)
  - [3.6.](#) Is the solution flexible enough to allow the participation of the end-nodes (CNs and/or MNs)? . . . . . [13](#)
  - [3.7.](#) Does the solution allow for a hierarchical scheme? . . . . . [13](#)
  - [3.8.](#) What is the target protocol complexity? . . . . . [14](#)
  - [3.9.](#) How is routing performed within the ATN? . . . . . [14](#)
  - [3.10.](#) Does the solution allow for implementing legal/political/economical requirements? . . . . . [14](#)
  - [3.11.](#) What is the robustness of the solution (i.e. what type of failure affects to the reachability)? . . . . . [15](#)
- [4.](#) Security Considerations . . . . . [15](#)
- [5.](#) IANA Considerations . . . . . [15](#)
- [6.](#) Acknowledgments . . . . . [15](#)
- [7.](#) References . . . . . [15](#)
  - [7.1.](#) Normative References . . . . . [15](#)
  - [7.2.](#) Informative References . . . . . [16](#)
- [Appendix A.](#) Change Log . . . . . [17](#)
- Authors' Addresses . . . . . [17](#)
- Intellectual Property and Copyright Statements . . . . . [18](#)

## **1. Introduction**

This document assumes that the reader is familiar with the terminology related to Network Mobility [4] and [5], and with the Mobile IPv6 [2] and NEMO Basic Support [3] protocols.

The MEXT WG is currently chartered to work on three use cases for route optimization of network mobility, namely aeronautics [6], vehicular [7] and consumer electronics [8]. The work on the requirements for the aeronautics use case seems to be mature enough at this point to start discussing about solutions. This document is an initial attempt aimed at fostering discussion on solutions, by presenting a general framework of how a solution for the aeronautics use case could look like, and identifying and highlighting relevant questions, issues and deployment models that need to be taken care of during the solution definition process.

The requirements for the aeronautics use case [6] differentiate among three different domains of interest: Air Traffic Services (ATS), Air Operational Services (AOS) and Passenger Information and Entertainment Services (PIES). Besides, two kind of requirements are identified: required (minimal properties that a solution must possess) and desirable (difficult to quantify or not immediately needed requirements) characteristics. Since the PIES domain is not critical to safety-of-life, but mostly involves added comfort and business services to passengers, this domain has not been taken into account as input for the required characteristics.

Due to the very different nature of the required and desirable characteristics, and the importance of the former ones, this document only analyzes how a solution for ATS/AOS would look like.

## **2. Solution Space analysis**

In this section we try to outline the general lines of a NEMO Route Optimization solution for the aeronautics use case (ATS/AOS domains), based on the set of requirements described in [6], which we summarize next:

- o Separability: an RO solution MUST support configuration by using a dynamic RO policy database, so RO for certain flows can be disabled/enabled. A granularity level similar to the one of IPsec security policy databases is expected to be supported.
- o Multihoming: an RO solution MUST support multi-interfaced MRs, and it MUST allow the use of different interfaces (and also different MNPs) for different domains.

- o Latency: an RO solution MUST allow packets to use the MRHA tunnel while setting up or reconfiguring the RO path.
- o Availability: an RO solution MUST NOT prevent to fall-back using the default MRHA tunnel if the RO path fails for whatever reason. This basically also means that an RO solution MUST NOT introduce any new single point of failure for the communications.
- o Packet Loss: an RO solution SHOULD NOT cause either additional loss or duplication of data packets due to the use of RO, above that caused in the NEMO basic default solution.
- o Scalability: an RO solution MUST be simultaneously usable by hundreds of thousands of crafts without overloading the ground network or the routing system.
- o Efficient Signaling: an RO solution MUST be efficient in terms of the number of required signaling messages, and avoid signaling storms as a result of providing multiple ongoing flows with RO following a handover.
- o Security: an RO solution MUST NOT expose MNPs on the wireless egress link, MUST allow the receiver of BUs to validate CoA ownership, and MUST ensure that only explicitly authorized MRNs are able to send a BU for a specific MNP.
- o Adaptability: an RO solution MUST NOT prevent applications from using transport protocols, IPsec or new IP options.
- o Although it is not explicitly listed as a required characteristic -- but only suggested in [6] --, it seems to be widely accepted that modifications to CNS MUST NOT be required by an RO solution.

From this list of requirements, a first conclusion that can be obtained is that a solution for the aeronautics NEMO RO use case MUST NOT require changes on the CNS in order to correctly operate, that is, a solution MUST provide certain level of RO with legacy IPv6 CNS. This means that the solution would likely rely on a set of entities at the infrastructure, performing the RO function between them or/and also the MR.

In a glimpse, a solution along the lines mentioned before would operate as follows (Figure 1): an optimized route between a mobile network deployed in a craft and a CN (or set of CNS) is set-up (upon some sort of trigger/policy) between two (or more in general terms) NEMO RO entities (NROEs). These RO entities should be located -- in order to provide an optimized route as shorter as possible -- close to the mobile network and/or the CN. It should be noted that one of these RO entities may be collocated within the MR. A tunnel between the RO entities (or a chain/nesting of tunnels, in case several RO entities are involved in the same optimized route) is established, so data traffic between the mobile network and the CN (or set of CNS) can be routed through this shorter route (compared with the default MRHA one). It might be the case that certain additional operations are needed in order to ensure that traffic sent/received by the CN

(or by the mobile network) is routed through the R0 entities, so they can forward packets using the optimized route. Involving more than two R0 entities might be useful in order to deploy hierarchical schemes (i.e. a chain of R0 entities is set up along the path between the MR and the CN), in which the MR would only need to update/exchange signaling with the closest R0 entity, but not with all the R0 entities involved in the optimization. This may improve the handover performance for the optimized flows and reduce the signaling load.

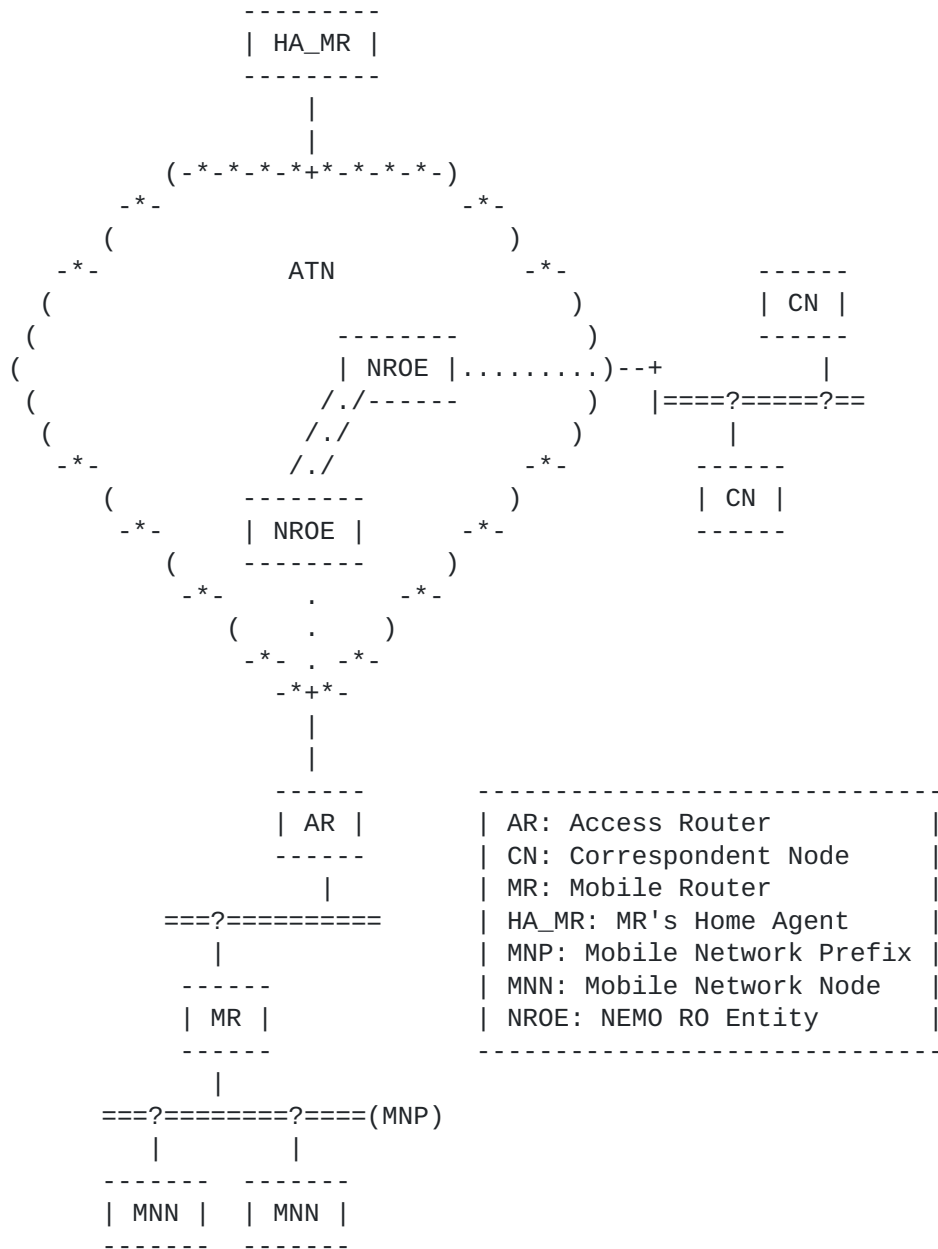


Figure 1: R0 entity-based solution architecture

This approach, based on the use of R0 network entities that are in charge of performing the NEMO R0, seems to be the solution that has received more positive feedback from the MEXT WG so far. This solution -- as described in this document -- is very general and leaves (on purpose) many aspects open/undefined. Depending on the particular design decisions that can be taken, completely different

solutions might be the outcome. For example, there are currently two proposed solutions that implement the RO entity-based concept in different ways: the global Home Agent to Home Agent (HAHA) [9], [10], and the Correspondent Router based RO for NEMO (CRON) [11].

Since different design decisions might result into completely different solutions, each of them meeting different requirements and providing different features, it is very important to understand the impact of the related design decisions, as well as the involved trade-offs, when designing a NEMO RO solution for the aeronautics use case. It is also important to look at the deployment issues derived from the particular characteristics of the Aeronautical Telecommunications Networks (ATNs) [12]. The next section of this document is aimed at identifying the different important design aspects, deployment issues and resulting trade-offs when considering an RO entity-based NEMO RO solution. The goal of such an exercise is to help the MEXT WG in the design of the NEMO RO solution for the aeronautics use case.

### **3. Design issues/questions/trade-offs**

In this section, we attempt to identify relevant design issues, questions and involved trade-offs when considering an RO entity-based NEMO RO solution, by asking the following questions:

1. Where are the RO entities located?
2. Who administratively manages the RO entities?
3. Which kind of addresses are gonna be used and who own them?
4. How many RO entities are needed to globally perform NEMO RO?
5. What trust relationships are needed?
6. Is the solution flexible enough to allow the participation of the end-nodes (CNS and/or MNNS)?
7. Does the solution allow for a hierarchical scheme?
8. What is the target protocol complexity?
9. How is routing performed within the ATN?
10. Does the solution allow for implementing legal/political/economical requirements?



11. What is the robustness of the solution (i.e. what type of failure affects to the reachability)?

### **3.1. Where are the RO entities located?**

A first important question is where the RO entities are located.

One option is to place the RO entities at the gACSPs. In this case the optimized path would have at least one end-point (depending on whether the solution involves two RO entities at the infrastructure or one entity at the infrastructure and the MR) at the gACSP. This may be not good enough from a performance point of view, since the farther the RO entities are from the communication end-points (i.e., the mobile network and the CNS) the more likely the resulting path would be less optimal.

There are some potentially relevant questions related to placing entities at the gACSPs:

- o can an MR get connectivity from two different gACSPs?
- o is it possible that the same MR needs to make use of RO entities placed at different gACSPs?
- o would it be possible/required that an RO path is set-up between RO entities belonging to two different gACSPs? If so, the different routing policies that gACSPs might implement are relevant, as we analyze later.

Another potential configuration is to place the RO entities at the ANSPs. This configuration would allow to have one end-point of the optimization close to the CNS (for the ATS scenario) and therefore would result in paths that in general would be closer to the optimal case. On the other hand, this approach would require more RO entities to be deployed, therefore eventually increasing the complexity of the solution. Besides, a solution that only places the infrastructure RO entities at the ANSPs would require the MR being the other RO entity in the NEMO RO process, since an MR might get connectivity through a gACSP, or through an lACSP that is not an ANSP. In the AOS scenario, this configuration might not work, since AOS CNS might not get connectivity through an ANSP, and therefore an RO entity should be deployed somewhere else.

In those communication scenarios in which the MR is attached to an ANSP access network and it is communicating with a CN also attached to the same ANSP, placing the RO entities at the ANSP (or collocated within the MR) provides the additional advantage that these communications would survive failures on the gACSPs to which the ANSP gets connectivity from. This brings the following question:

- o if a craft attached to an ANSP access network is communicating with a CN attached to the same ANSP, is it required for the NEMO R0 solution to survive when the link of the ANSP to its gACSP goes down? or put in a different way, would it be OK for such a communication to be broken? It should be noted that the default MRHA path used by the NEMO Basic Support protocol would likely fail in this scenario.

Another approach is to deploy infrastructure R0 entities at the networks where CN are attached (these networks might be ANSPs in some scenarios) and at MRs. This would provide shorter paths, at the cost of higher complexity.

Last, a solution might not assume any particular placement of the R0 entities, i.e. they can be located anywhere. This assumption, however, might not hold, depending on different aspects -- such as security and addressing (for example if prefixes used in ATS cannot leak to the Internet, leading to ATS traffic traversing the public Internet).

### **3.2. Who administratively manages the R0 entities?**

It is also important to analyze who will be the stakeholders than manage the R0 entities, since this might have a critical impact on the trust relationships that can be assumed among the different R0 entities.

One first scenario is the one in which all the R0 entities are managed by the same administrative entity. This compresses both the case of R0 entities deployed and managed by the airline company, and the case of a global ACSP providing the R0 entities for airlines with an agreement with the ACSP. The obvious advantage of this scenario is that it makes security and authentication easier, since all the R0 entities belong to the same administrative domain. However, this does not mean that this scenario is excluded from having trust issues, since a particular solution might require to inject routes in some parts of the network (e.g., R0 entities owned by an airline and placed in networks not managed by the airline, anycast routing, etc.), and this could require additional trust relationships.

A second scenario is that in which R0 entities are managed by different administrative domains. This approach has the advantage that it provides additional flexibility, but it has the drawback of requiring additional trust agreements in order to enable NEMO R0 to be provided in a secure way. Depending on the particular solution, these additional trust relationships may include those that are necessary to enable anycast routing, route injection or strong state synchronization, among others. These are examples of functions that

are usually not easy to achieve across different domains.

For AOS, it seems assumable to deploy an RO entity close the the CNs, and then perform RO between this entity and the MR, since both entities are operated by the same organization (therefore, existence of certificates between these nodes could be expected). The ATS case is different, and should be analyzed carefully. Although some trust may exist between an RO entity belonging to an ANSP and another RO entity (e.g., one deployed at the aircraft, or one deployed at one gACSP), assuming the existence of certificates or strong trust relationships is not clear at this point [12]. This brings the following question:

- o which trust relationships are expected? this will be analyzed in further detail in another subsection.

### **3.3. Which kind of addresses are gonna be used and who own them?**

Addressing aspects might be relevant for the design of a NEMO RO solution. Some related questions are the following:

- o are there gonna be reserved blocks of addresses for aeronautical use (i.e. addressing used to derive MNPs from)?
- o can prefixes used to derive the MNPs configured at the crafts leak on the routing tables of the public Internet? can ATS/AOS traffic traverse the public Internet?
- o what kind of addressing is gonna be used for ATS and AOS? would it be the same kind of addressing?
- o is it fine to use PA addresses to derive the MNPs configured at the crafts or is it a requirement to use PI addresses (delegated to the airline, to enable provider independency)? One possible solution design is that MNPs are derived from the addressing of a gACSP which deploy several RO entities to perform NEMO RO, but this scenario would tie the airline to keep using the same provider.

### **3.4. How many RO entities are needed to globally perform NEMO RO?**

Another important aspect that should be taken into account is the number of RO entities that would be required to perform NEMO RO efficiently. There are many aspects that may have an impact on the number of required RO entities, such as:

- o Location of the RO entities. If a solution is based on placing RO entities very close to the CNs this might require an entity per CN network (e.g., per ANSP). Solutions relying on RO entities located at gACSPs may require less entities to be deployed.
- o Who owns/manages the RO entities. Depending on the particular solution, it could happen that each airline has to deploy its own RO entities, thus requiring a set of RO entities per airline.

- o Required level of RO. Of course, depending on the optimization levels that are required to be achieved, the location and number of RO entities would change. If certain amount of additional delay are allowed, it is expected that less entities would be needed, since there is usually a trade-off between the number and location of RO entities, and the reduction of the delay due to the optimized path.

### **3.5. What trust relationships are needed?**

Trust relationships are a quite important aspect to be analyzed. As it has been described in this document, a solution based on the deployment of RO entities may take many different forms, depending on the design decisions and the deployment assumptions that are followed. Most of the design decisions would have an impact or would be constrained by the trust relationships that are in place among the different players involved in the NEMO RO.

A solution based on the establishment of an optimized path between two or more RO entities inherently requires those entities to have strong trust relationships with the end-points of the communications, since they are providing an alternative -- over the MRHA default path -- route for their communication. Therefore, both MNNs and CNS MUST have some form of trust relationship with the RO entities, to allow the latter set-up an optimized route for their traffic (on their behalf). As an example, both the MR and the HA of a particular mobile network clearly have the required trust relationship with the MNNs of the mobile network, and therefore they could take part of an optimization mechanism. Other entities but the MR and its HA would need additional trust relationships in place in order to take part of a NEMO RO solution.

The RO entities involved in a NEMO RO solution MUST also have some trust relationship between them, allowing them to authenticate each other.

Additionally, RO entities involved in an RO attempt MUST be able to show each other that they are authorized to send and receive packets originated/destined to the nodes (MNNs or CNS) they are providing RO with. As an example, let's assume a particular solution in which there are two RO entities, one placed close to the mobile network, and the other placed close to the CN. In this example, the entity close to the mobile network should be able to show to the one close to the CN that it is authorized to send/received packets originated/destined to the MNNs of that particular mobile network. It should be noted that the same kind of authorization is required and provided when the NEMO Basic Support protocol is used (i.e. the MR has to be authorized to set-up a tunnel with its HA to exchange packets, and

both MR and HA have to authenticate each other before setting up the tunnel). Actually, the same authorization is required between any Internet host and the routers it uses to forward its traffic.

RO entities MUST also be authorized to inject the routes (if any) required to get the packets that are subject of being route optimized. The simpler case is that in which an RO entity is the default router of the MNs (i.e. the MR) or the CNs, since in this scenario nothing is required to make MNs/CNs forward to the RO entity their traffic, and therefore this entity is inherently authorized to forward that traffic. Other kind of solutions, in which the RO entities are not collocated with the MR and the default router of the CN might require the RO entities to inject routes within a certain portion of the network. This might be hard to achieve across different domains.

The location and ownership of the RO entities would likely have a great impact on the potentially required trust relationships. Therefore, trust and location issues have to be simultaneously considered.

**3.6. Is the solution flexible enough to allow the participation of the end-nodes (CNs and/or MNs)?**

Supporting legacy end-nodes (MNs and CNs) seems to be a required characteristic, although it is not explicitly listed as that in [6]. That means that a solution MUST NOT require changes neither at the MNs nor at the CNs to operate. However, that does necessary imply that a particular solution cannot benefit from inserting changes on some specific MNs and/or CNs. In other words, a solution could provide the option of collocating the RO entity function within some MNs and/or CNs -- in those scenarios in which these modifications can be done. This brings the following question:

- o is it permitted for a solution to collocate the RO entity function within certain MNs and/or CNs in case their software upgrade is possible and that change brings operation benefits?

**3.7. Does the solution allow for a hierarchical scheme?**

Solutions based on the deployment of RO entities that perform the required route optimization operations may benefit from adopting hierarchical schemes. This, for example, may help to reduce signaling and produce faster handovers. Therefore, a consideration that could also be taken into account when designing a solution is if it would support a hierarchical mode of operation.

Another somehow related design consideration is the following: a particular solution might benefit from deploying NetLMM-alike access

networks and collocating the functionality of the NEMO RO entity with that of the LMA. This could improve the overall performance, although at the prize of increasing the global complexity and requiring ACSPs to be NetLMM-alike.

### **3.8. What is the target protocol complexity?**

It is obvious that a solution should be as less complex as possible, but there is always a trade-off involved: less complex solutions usually provide less features/performance gains/etc., and the other way around. There are some particular requirements of the aeronautical NEMO RO scenario that would likely impact on the solution complexity, and that should be taken into account.

For example, in order to meet the separability requirement [12], RO entities in charge of performing the RO would have to be able to decide whether a certain flow has to be optimized or not. This could be done by local policies or explicit signaling. Even in the case of local policies, some mechanisms would be needed to support the update/modification of the policies. It seems likely that it would be up to the mobile networks to decide what flows are to be optimized and which not. Therefore, the costs associated to meet the separability requirement would likely involve some sort of signaling between the mobile networks and the RO entities (at least to trigger the NEMO RO of a particular flow). This cost should be evaluated and taken into account when designing an RO entity-based solution. As an example, if a solution collocates one RO entity function within the MR, this solution would likely require less signaling to meet the separability requirement than another solution that makes use of RO entities placed on the network infrastructure.

### **3.9. How is routing performed within the ATN?**

Routing policies and related issues within the ATN are an important input to be considered when designing an RO entity-based solution. Therefore, we should address the following questions:

- o is it OK to have asymmetrical optimized routes? depending on the design of the solution, it might be possible that some sort of asymmetric routing appears.
- o what are the routing policies followed by the ACSPs (especially gACSPs)? do they do cold or hot-potato routing? cold-potato routing may lead to very suboptimal routes under some particular scenarios, even when a NEMO RO solution is used.

### **3.10. Does the solution allow for implementing legal/political/economical requirements?**

In some Internet scenarios, it is preferred that data traffic does

not traverse certain networks because of different reasons, such as legal, political or economical ones. Is that also the case for the aeronautics use case?. If so, it might be important to provide NEMO RO solutions with the required mechanisms to implement the policies that translate those potential legal/political/economical requirements.

**3.11. What is the robustness of the solution (i.e. what type of failure affects to the reachability)?**

It is also important to analyze the robustness of a particular solution design, in terms of the types of failures that might affect to the reachability of the network. For example, a solution may provide an RO path despite of a broken path between the NEMO and its home network, while another one may not. It is important to identify which are the failures that can happen in an ATN, which ones would only affect the reachability of a craft only when using a NEMO RO solution, and if it is fine to have those failures.

**4. Security Considerations**

This document analyzes a general approach to perform NEMO RO for the aeronautics use case. As such, it identifies some security issues that should be taken into account in the design of a concrete solution.

**5. IANA Considerations**

This document has no actions for IANA.

**6. Acknowledgments**

The work of Carlos J. Bernardos has been partly supported by the Spanish Government under the POSEIDON (TSI2006-12507-C03-01) project.

**7. References**

**7.1. Normative References**

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [2] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), 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] Manner, J. and M. Kojo, "Mobility Related Terminology", [RFC 3753](#), June 2004.
- [5] Ernst, T. and H-Y. Lach, "Network Mobility Support Terminology", [RFC 4885](#), July 2007.
- [6] Eddy, W., Ivancic, W., and T. Davis, "NEMO Route Optimization Requirements for Operational Use in Aeronautics and Space Exploration Mobile Networks", [draft-ietf-mext-aero-reqs-02](#) (work in progress), May 2008.
- [7] Baldessari, R., Ernst, T., Festag, A., and M. Lenardi, "Automotive Industry Requirements for NEMO Route Optimization", [draft-ietf-mext-nemo-ro-automotive-req-01](#) (work in progress), July 2008.
- [8] Ng, C., Hirano, J., Petrescu, A., and E. Paik, "Consumer Electronics Requirements for Network Mobility Route Optimization", [draft-ng-nemo-ce-req-02](#) (work in progress), February 2008.
- [9] Thubert, P., Wakikawa, R., and V. Devarapalli, "Global HA to HA protocol", [draft-thubert-mext-global-haha-00](#) (work in progress), March 2008.
- [10] Wakikawa, R., Shima, K., and N. Shigechika, "The Global HAHA Operation at the Interop Tokyo 2008", [draft-wakikawa-mext-haha-interop2008-00](#) (work in progress), July 2008.
- [11] Bernardos, C., Calderon, M., and I. Soto, "Correspondent Router based Route Optimisation for NEMO (CRON)", [draft-bernardos-mext-nemo-ro-cr-00](#) (work in progress), July 2008.
- [12] Bauer, C. and S. Ayaz, "ATN Topology Considerations for Aeronautical NEMO R0", [draft-bauer-mext-aero-topology-00](#) (work in progress), July 2008.



## Appendix A.    **Change Log**

Changes from -00 to -01:

- o Terminology changes: s/correspondent entity/R0 entity.
- o New solution design issue: robustness.
- o Marcelo Bagnulo enlisted as author.
- o Some editorial changes.

### Authors' Addresses

Carlos J. Bernardos  
Universidad Carlos III de Madrid  
Av. Universidad, 30  
Leganes, Madrid 28911  
Spain

Phone: +34 91624 6236  
Email: [cjbc@it.uc3m.es](mailto:cjbc@it.uc3m.es)  
URI:    <http://www.it.uc3m.es/cjbc/>

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

Phone: +34 91624 9500  
Email: [marcelo@it.uc3m.es](mailto:marcelo@it.uc3m.es)  
URI:    <http://www.it.uc3m.es/marcelo/>

## Full Copyright Statement

Copyright (C) The IETF Trust (2008).

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](mailto:ietf-ipr@ietf.org).