

MEXT Working Group
Internet-Draft
Intended status: Informational
Expires: March 13, 2010

C. Bauer
S. Ayaz
DLR
A. Ebalard
EADS
September 9, 2009

Solution Space for Aeronautical NEMO RO
draft-bauer-mext-aero-solSPACE-00

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [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/lid-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 March 13, 2010.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Internet-Draft

Aeronautical NEMO RO Sol. Space

September 2009

Abstract

Many potential solutions have been proposed for NEMO Route Optimization, although none has been adopted up to now. This draft aims on evaluating the different approaches for the aeronautical use case. At the end, a recommendation for the next steps is given.

Table of Contents

1.	Introduction	3
2.	Overview	4
3.	Discussion	6
3.1.	Correspondent Router	6
3.2.	Global HA to HA	9
4.	Next steps	12
5.	Considerations for PIES	13
6.	Security Considerations	14
7.	Acknowledgements	15
8.	References	16
8.1.	Normative References	16
8.2.	Informative References	16
Appendix A.	Short Overview of all RO Solutions	19
A.1.	Analysis	19
A.2.	Applicability to the Aeronautical Environment	27
	Authors' Addresses	29

1. Introduction

An extensive overview of NEMO Route Optimization (RO) solution candidates has been provided in [[RFC4889](#)]. The options are manifold and obviously the solution has to be adopted based on the operational environment. The task of this draft is to investigate the solutions in more detail and highlight the deficiencies of the protocols that should be addressed in the future.

A document listing the requirements that have to be fulfilled by a potential aeronautical NEMO RO solution has already been compiled [[I-D.ietf-mext-aero-reqs](#)]. In addition, another document [[I-D.bauer-mext-aero-topology](#)] provides information on the aeronautical network environment. We will rely on both to perform the investigation.

It is expected that the reader is familiar with the NEMO Support Terminology [[RFC4885](#)] and the three above mentioned documents.

2. Overview

The focus of our investigations is on R0 for the ATS/AOS service classes (cf. [[I-D.ietf-mext-aero-reqs](#)]). As defined in [[I-D.ietf-mext-aero-reqs](#)], the problem of nested MRs is not a requirement and merely desired and can therefore be ignored in the first step. Its only use case is that of a MANET of aircraft, where other solutions to the problem can be found, e.g. within the MANET (routing) itself, that are transparent to NEMO. Support for VMNs is also not required for the ATS/AOS domain.

Hence the applicable NEMO R0 solutions can be categorized as follows, listing those nodes that are involved in mobility related signaling:

1. MNN to CN: R0 is performed between the end systems.
2. MR to CN: The Mobile Router performs R0 on behalf of the MNN with the CN.
3. MR to CR: The MR performs R0 with a Correspondent Router that is located close to the CN and that can forward traffic from and to the CN.
4. MR to HA: The MR binds to a topologically closer Home Agent.

The requirements from [[I-D.ietf-mext-aero-reqs](#)] play an important role for the analysis. We list them here for ease of reading:

1. Req1 - Separability: "Since R0 may be inappropriate for some

flows, an RO scheme MUST support configuration by a per-domain dynamic RO policy database [...]. The rationale for this requirement is to trigger RO only for flows that really require it.

2. Req2 - Multihoming: " An RO solution MUST support an MR having multiple interfaces, and MUST allow a given domain to be bound to a specific interface. It MUST be possible to use different MNPs for different domains". Multihoming itself is supposed to allow segregation of traffic flows over different interfaces.
3. Req3 - Latency: "While an RO solution is in the process of setting up or reconfiguring, packets of specified flows MUST be capable of using the MRHA tunnel".
4. Req4 - Availability: "An RO solution MUST be compatible with network redundancy mechanisms and MUST NOT prevent fall-back to the MRHA tunnel if an element in an optimized path fails. An RO mechanism MUST NOT add any new single point of failure for

communications in general". Redundancy mechanisms do not have to be considered for the RO mechanism itself; this requirement merely tries to ensure that RO does not interfere with any existing redundancy mechanism.

5. Req5 - Packet Loss: "An RO scheme SHOULD NOT cause either loss or duplication of data packets during RO path establishment, usage, or transition, above that caused in the NEMO basic support case. An RO scheme MUST NOT itself create non-transient losses and duplications within a packet stream".
6. Req6 - Scalability: "An RO scheme MUST be simultaneously usable by the MNNs on hundreds of thousands of craft without overloading the ground network or routing system. This explicitly forbids injection of BGP routes into the global Internet for purposes of RO".
7. Req7 - Efficient Signaling: "An RO scheme MUST be capable of efficient signaling in terms of both size and number of individual signaling messages and the ensemble of signaling messages that may simultaneously be triggered by concurrent flows".

8. Req8 - Security #1 (for ATS): "The RO scheme MUST NOT further expose MNPs on the wireless link than already is the case for NEMO basic support".
9. Req8 - Security #2 (for ATS): "The RO scheme MUST permit the receiver of a BU to validate an MR's ownership of the CoAs claimed by an MR" and "the RO scheme MUST ensure that only explicitly authorized MRs are able to perform a binding update for a specific MNP".
10. Req9 - Adaptability: "Applications using new transport protocols, IPsec, or new IP options MUST be possible within an RO scheme". In particular, the RO scheme should not fail on the use of previously unknown higher layer protocols.

In the first stage we have performed a short analysis of all the four possible approaches outlined above. We came to the conclusion that only two of these, namely "MR to CR" and "MR to HA", are of interest to the aeronautical environment. We therefore focus on these in [Section 3](#) below. Nevertheless, for completeness, we have added the short analysis of all four categories in [Appendix A](#).

[3.](#) Discussion

This section discusses the "MR to CR" and "MR to HA" proposals in detail.

[3.1.](#) Correspondent Router

The CR approach was first introduced in [[I-D.orc](#)] and later a proposal specifically targeted at the aeronautical requirements has been specified in [[I-D.draft-bernardos-mext-nemo-ro-cr](#)]. Another document [[I-D.draft-wakikawa-mext-cr-consideration](#)] discusses possible issues and provides considerations for the CR protocol.

Req1 - Separability: at the moment, the only obvious way how separation could work is on a per address/prefix basis and not

traffic type or application. The major problem is related to the CR relying on IGP advertisements within its network to attract traffic destined for the MR. In that case, the separability of different flows can only be achieved based on addresses/prefixes, but not on traffic type on the CR side. The CR specifications (either [[I-D.orc](#)] or [[I-D.draft-bernardos-mext-nemo-ro-cr](#)]) do not discuss the compatibility with [[I-D.draft-ietf-monami6-multiplecoa](#)] and [[I-D.draft-ietf-mext-flow-binding](#)] in order to dispatch flows between the MN and its CR via different interfaces (CoA) on the MR. However, whether it is really necessary to direct certain flows, destined to the same CN, over the R0 path and others over the MR-HA tunnel is not obvious. At least for the ATS domain, it could be considered sufficient to have separability on a per-address basis.

Req2 - Multihoming has, to a certain degree, already been addressed above. Nevertheless, when taking a more detailed look at MCoA, its implementation within the CR protocol comes at the expense of the MR having to register each CoA separately (no bulk registration). The signaling overhead for the registration grows linearly with the number of CRs the MR is binding with. At least for the ATS domain, this might not be considered as problematic given that only one or at most two CR bindings are probably active at a certain point in time (cf. [[I-D.bauer-mext-aero-topology](#)]).

Req3 - Latency: the delay of establishing R0 with the CR is the sum of the discovery and the registration procedures. [Appendix A.3](#) of [[I-D.draft-bernardos-mext-nemo-ro-cr](#)] states that these steps may be done while the MR-HA tunnel is still used for sending user data. The signaling during a handover is equivalent to that of standard MIPv6 R0 for [[I-D.orc](#)] and an IKEv2 and BU/BA exchange for [[I-D.draft-bernardos-mext-nemo-ro-cr](#)]. For as long as this signaling is performed, it could be expected to rely on the (updated) MR-HA tunnel for as long as the binding with the CR has not been renewed.

This behaviour depends though on which of the MR-CR and MR-HA bindings has been first successfully updated. A change of the CoA will trigger the same series of events, except for the discovery procedure that is not needed anymore.

Req4 - Availability: the MR maintains a communication path with a unique and static HA and also has direct routing paths (R0 paths) with its correspondents, via the various CRs located near those final

correspondents. In case of failure of the CR, the MR could switch back to the MR-HA tunnel, given that the failure is detected. However, it is difficult to assess whether the probability for a CR failure is larger than for a HA failure. On the other hand, a CR could also increase the availability in case of HA failure if credentials between MR and CR are available and RO can therefore be established in a different way compared to MIPv6 RO. In terms of routing path characteristics, the shorter path used by the MR-CR tunnel might have a lower probability of failure than the "longer" MR-HA path.

Req5 - Packet loss: as already mentioned for Req3, while performing signaling for setting up the RO state, the MR-HA tunnel could be used for user data and no packet loss would therefore be induced. For as long as the mobility binding at the CR has not been re-established, packets can flow from the MR via the HA (if already available) directly to the CN, but on the reverse path packets will be sent to the old CoA of the MR and therefore be lost.

Req6 - Scalability: if the number of CNs and, most importantly, the number of networks in which they are hosted increases, more CRs will be needed. [Appendix A.6](#) of [I-D.[draft-bernardos-mext-nemo-ro-cr](#)] does not discuss the fact that the more CRs a MR has a binding with, the more signaling it needs to send after a handover (change of CoA). It is therefore important to consider the number of nodes within the ATS and AOS domains. Taking into account the ATS communication model (cf. [I-D.[bauer-mext-aero-topology](#)]), the number of CRs that the MR should have an active binding with might be limited to two.

Req7 - Efficient Signaling: signaling is required for each binding with a CR. If MCoA/flow-binding should be supported, the amount of signaling might increase. An aspect that is not discussed in [Appendix A.7](#) of [I-D.[draft-bernardos-mext-nemo-ro-cr](#)] is the signaling that may be required for the detection of the CR, required once at the time RO is initiated.

Req8 - Security: a main requirement is that the authenticity of the MNP of the MR has to be verified. [I-D.[orc](#)] relies on the well known Return Routability method of Mobile IPv6 and therefore inherits its security weaknesses. [I-D.[draft-bernardos-mext-nemo-ro-cr](#)] proposes

to use IKEv2/IPsec to secure the BU/BA exchange with the assumption

that "certificates are used to prove ownership of prefixes by MRs and CRs". The location, domain and associated authority under which the CR is deployed can make the deployment difficult, depending on the availability of those credentials.

Req9 - Adaptability: for as long as generic packet tunneling is used between MR and CR, no problems can be expected wrt security protocols within the inner tunnel. This is similar to NEMO Basic Support. The discussion related to Req1 and Req2 already mentioned that using certain traffic types (e.g. specific transport protocol) for either flow-specific RO or flow-bindings management could be regarded as problematic.

The following paragraphs are not directly related to the requirements anymore, but are from a more general perspective.

A critical item, that is not directly covered by the requirements, is the discovery of the CR: it is difficult for the MR to know the anycast address for a particular CR, as needed for the discovery request message. Deriving it from the CN address, as originally proposed in [[I-D.orc](#)], can not work if both CN and CR do not share the same 64bit prefix.

A CR could be regarded as an entity that is close to the CN but already has a trust relationship with the MR. Taking a closer look at the topology of the aeronautical environment, presented in [[I-D.bauer-mext-aero-topology](#)], reveals that the following two options for deployment of CRs are possible for ATS:

1. CR in ANSP network: while the topological location is excellent for RO purposes, the CR can not be considered to have a trust relationship anymore. If the CR would have, then a trust-relationship would also be possible with the CN, as both are within the same operational domain.
2. CR in the neighbouring gACSP network: having a trust relationship between MR and the gACSP domain is very likely, but the topological location is not ideal anymore. The CR can, in general, not be on-path anymore. In fact, it is located in a completely different domain, a situation that will cause operational problems, especially if the network of the CN is multi-homed and peering with several gACSPs.

These deployment options are dependent on the credentials mentioned in the discussion of Req8 above.

3.2. Global HA to HA

The Global Home Agent to Home Agent protocol is specified in [I-D.[draft-wakikawa-mext-global-haha-spec](#)] and makes use of the HA reliability protocol [I-D.[draft-ietf-mip6-hareliability](#)].

Req1 - Separability: in Global HAHA, instead of having RO triggered on a per-flow or per-destination basis, the MR/MN's position dictates the HA instance used by the MN/MR (usually the topologically closest one). From this perspective, HAHA does not directly provide a way to support this requirement. Whether this requirement is applicable to HAHA is a different question though. Nevertheless, while not explicitly discussed in HAHA at the time of writing, compatibility with MCoA [I-D.[draft-ietf-monami6-multiplecoa](#)] and flow-binding [I-D.[draft-ietf-mext-flow-binding](#)] specifications could address this issue.

Req2 - Multihoming: compatibility with the MCoA/flow-binding specifications needs to be discussed, as already shortly mentioned above in Req1. From a more general point of view, Global HAHA does not fundamentally modify the relationship with the HA, therefore making these protocol extension applicable to it. But due to the way the primary HA of a MN/MR is selected (the topologically closest instance of the first active CoA is used), the addition of new CoA raises questions. If two CoAs have different HA attractors, how should the protocol behave: should it manage to keep a single primary HA for all CoAs (considering one is a primary) or be extended to support bindings for different CoAs with different HAs? Because the binding cache has to be shared between the HA instances, adding support for MCoA/flow-binding to Global HAHA may require additional synchronization and complicates the protocol.

Req3 - Latency: the overall latency is composed of the BU/BA exchange between the MR and the HA. In case the MR moves to a different location that attracts a different HA, the MR will receive a HA Switch Message that forces him to establish a binding with the new HA. This requires an additional IKEv2/IPsec and BU/BA exchange. In general, the latency is therefore equivalent to the standard NEMO protocol and several additional RTTs if the MR binds to a new HA after handover events (change of CoA).

Req4 - Availability: the operation of HAHA is based on the HA reliability protocol [I-D.[draft-ietf-mip6-hareliability](#)] that describes how failovers are performed between local instances. Due to this, a local HA failure can be overcome and this part of the requirement be therefore fulfilled due to the additional local HA

instances. When looking at the MR-HA path, if the routing to the home network is broken (not the physical link the MR is using but an

element along the path), current behavior is unknown (both for Global HAHA and NEMO Basic Support). Depending on the precise nature of the routing failure, it might be possible within the context of Global HAHA that another HA island starts attracting traffic (this however will also be dependent on the BGP convergence time of the anycast prefix advertised by the "new" HA island).

Req5 - Packet loss: Global HAHA shows the same behaviour as standard MIPv6/NEMO Basic Support. For as long as the MN/MR has not finished mobility signaling for updating the tunnel with the new CoA, the HA will send packets to the old CoA that will therefore be lost. After this signaling has been finished, the MR could be informed by its current HA to switch to another, closer HA. While this adds latency for setting up the RO path (shorter route due to closer HA), it is a "soft" movement and will not cause packet loss as long as the binding with the old HA is kept active while the new binding has not yet been successfully established.

Req6 - Scalability: with Global HAHA, a given MR being handled by the closest HA instance leads to a natural distribution of the traffic between all the HAs. The traffic load at the ground network is dependent on the location of the MRs and the HA islands. From a routing table perspective, Global HAHA deployment puts some load on the BGP routing system, as prefixes have to be advertised. The number of BGP advertised prefixes is dependent on the number of anycast prefixes and HA islands, but should at least be constant and not show any frequent advertise/withdrawal behaviour. Another critical aspect is the network traffic caused by synchronizing the binding caches between the various HA instances (cf. [I-D.[draft-ietf-mip6-hareliability](#)]) if the number of HAs is very large.

Req7 - Efficient Signaling: in Global HAHA the amount of signaling between the MR and its HA is roughly the same as in NEMO Basic Support. This number is increased by the HA switch operation and therefore depending on the number of HAs. Most importantly, the amount of signaling is constant - it does not depend on the number of correspondent nodes the MR is currently communicating with. Considering that MCoA/flow-bindings will work with Global HAHA in the

future, the expected amount of signaling will still remain minimal, for as long as it is performed with only a single HA. As already discussed for Req2 - Multihoming, synchronizing flow-bindings among different HAs will induce more overhead.

Req8 - Security: as in Global HAHA bindings are only performed with the HAs of the home network (within a single administrative domain), meeting the security requirements is much easier to fulfill. IKEv2/IPsec protects the BU/BA message exchanges between the MR and the HA

instances. Consecutively, due to the use of IKEv2, the problem of MNP authentication is directly addressed.

Req9 - Adaptability: As Global HAHA does not modify the operation of the MR-HA tunnel (apart from having it with the closest HA instance), no limitations are introduced when compared to NEMO Basic Support.

The following paragraphs are not directly related to the requirements anymore, but are from a more general perspective.

Global HA to HA does not have the problems of the CR protocol wrt security and credentials, as its scope is restricted to the MR and the Mobility Service Provider (MSP) operating the HAs. Its disadvantage is that the level of provided RO depends on the global network presence of the MSP. gACSPs (cf.

[\[I-D.bauer-mext-aero-topology\]](#)) can achieve this goal, whether this also holds for airlines acting as their own MSP is unclear. This is in fact a critical aspect for HAHA, as a home network without a distributed global presence can not meet the original goal of providing an adequate, short latency RO path.

As mentioned in Section 3.3.1 of [\[I-D.bauer-mext-aero-topology\]](#), the MR could actually be attached to the same network as the CN. Failures within the home network or with the routing path between the visited and the home network breaks connectivity to the HAs and would then prevent communications although both nodes (MR and CN) are in close topological proximity.

[4.](#) Next steps

The investigation showed that both protocols have advantages and disadvantages, although open questions remain for both of them.

The CR protocol has issues related to its discovery procedure, but can support multihoming, although at the expense of having to register each CoA separately. The deployment options are restricted by the availability of credentials and by the location of the CNs.

Global HA to HA has the advantage of restricting the mobility signaling to the MR and home network only. How multihoming can be addressed is not yet clear; in addition, the overhead caused in the ground network might become an issue with a growing number of HAs.

We therefore think that additional work and investigations in the following areas are required:

For CR:

1. Discovery procedure.
2. A more detailed specification of the mobility signaling between MR and CR, including mutual authentication between MR and CR.

For HAHA:

1. Overhead caused by binding cache synchronizations between HAs.
2. Support for multihoming as in [I-D.[draft-ietf-monami6-multiplecoa](#)] and [I-D.[draft-ietf-mext-flow-binding](#)].

Aftwards, the solution space investigation can be revisited and the proper solution be selected.

5. Considerations for PIES

All discussion up to now have been focussed on the ATS/AOS traffic classes. Passenger communications, especially for passenger owned devices, has been ignored.

We think though that the number of options is severely limited for this scenario:

1. RO signaling involving the CNs is unrealistic as those nodes are usually within the public Internet and will most like not implement any mobility functionality.
2. Deployment of special infrastructure in the Internet, e.g. CRs, just for the purpose of RO seems unlikely.
3. Forcing passengers to install software for mobility functionality might be regarded as problematic.

As a consequence, Route Optimization has to be limited to the MR and the MSP network. This implies Global HA to HA is a reasonable solution for the PIES domain. This would also allow to reuse the protocol for the ATS/AOS environment.

[6.](#) Security Considerations

This document only presents information related to the aeronautical NEMO RO solution space. There are no security issues in this document.

[7.](#) Acknowledgements

Christian Bauer and Serkan Ayaz have been partially supported by the European Commission through the NEWSKY project. The views and

conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the NEWSKY project or the European Commission.

8. References

8.1. Normative References

[I-D.bauer-mext-aero-topology]

Bauer, C. and S. Ayaz, "ATN Topology Considerations for Aeronautical NEMO RO", [draft-bauer-mext-aero-topology-00](#) (work in progress), July 2008.

[I-D.ietf-mext-aero-reqs]

Eddy, W., Ivancic, W., and T. Davis, "Network Mobility Route Optimization Requirements for Operational Use in Aeronautics and Space Exploration Mobile Networks", [draft-ietf-mext-aero-reqs-04](#) (work in progress), August 2009.

[RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.

[RFC3963] Devarapalli, V., Wakikawa, R., Petrescu, A., and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", [RFC 3963](#), January 2005.

[RFC4889] Ng, C., Zhao, F., Watari, M., and P. Thubert, "Network Mobility Route Optimization Solution Space Analysis", [RFC 4889](#), July 2007.

8.2. Informative References

[I-D.bernardos-nemo-miron]

Bernardos, C., "Mobile IPv6 Route Optimisation for Network Mobility (MIRON)", [draft-bernardos-nemo-miron-01](#) (work in progress), July 2007.

[I-D.[draft-bernardos-mext-nemo-ro-cr](#)]

Bernardos, C., Calderon, M., and I. Soto, "Correspondent Router based Route Optimisation for NEMO (CRON)", July 2008, <<http://tools.ietf.org/html/draft-bernardos-mext-nemo-ro-cr>>.

[I-D.[draft-ietf-mext-flow-binding](#)]

Soliman, H., Tsirtsis, G., Montavont, N., Giaretta, G., and K. Kuladinithi, "Flow Bindings in Mobile IPv6 and NEMO Basic Support", July 2009, <[draft-ietf-mext-flow-binding](#)><http://draft-ietf-mext-flow-binding>>I-D.[draft-ietf-mext-flow-binding](#)>.

[I-D.[draft-ietf-mip6-hareliability](#)]

Internet-Draft

Aeronautical NEMO RO Sol. Space

September 2009

July 2009,

<<http://tools.ietf.org/html/draft-ietf-mip6-hareliability>>.

[I-D.[draft-ietf-monami6-multiplecoa](#)]

Wakikawa, R., Tsirtsis, G., Ernst, T., and K. Nagami, "Multiple Care-of Addresses Registration", May 2009, <<http://tools.ietf.org/html/draft-ietf-monami6-multiplecoa>>.

[I-D.[draft-wakikawa-mext-cr-consideration](#)]

Wakikawa, R., "The Design Consideration of Correspondent Router", July 2008, <<http://tools.ietf.org/html/draft-wakikawa-mext-cr-consideration>>.

[I-D.[draft-wakikawa-mext-global-haha-spec](#)]

Wakikawa, R., Zhu, Z., and L. Zhang, "Global HA to HA Protocol Specification", July 2009, <<http://tools.ietf.org/html/draft-wakikawa-mext-global-haha-spec>>.

[I-D.ndproxy]

Lee, K., Jeong, J., Park, J., and H. Kim, "ND-Proxy based Route and DNS Optimizations for Mobile Nodes in Mobile Network", February 2004, <<http://tools.ietf.org/html/draft-jeong-nemo-ro-ndproxy>>.

[I-D.orc]

Wakikawa, R. and M. Watari, "Optimized Route Cache Protocol (ORC)", October 2004, <<http://tools.ietf.org/html/draft-wakikawa-nemo-orc>>.

[I-D.pd]

Lee, K., Jeong, J., Park, J., and H. Kim, "Route Optimization for Mobile Nodes in Mobile Network based on Prefix Delegation", February 2004, <<http://tools.ietf.org/html/draft-leekj-nemo-ro-pd>>.

[RFC3971]

Arkko, J., Kempf, J., Zill, B., and P. Nikander, "SEcure Neighbor Discovery (SEND)", [RFC 3971](#), March 2005.

[RFC4225]

Nikander, P., Arkko, J., Aura, T., Montenegro, G., and E. Nordmark, "Mobile IP Version 6 Route Optimization Security Design Background", [RFC 4225](#), December 2005.

- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", [RFC 4862](#), September 2007.
- [RFC4885] Ernst, T. and H-Y. Lach, "Network Mobility Support Terminology", [RFC 4885](#), July 2007.
- [RFC5142] Haley, B., Devarapalli, V., Deng, H., and J. Kempf,

Bauer, et al.

Expires March 13, 2010

[Page 17]

Internet-Draft

Aeronautical NEMO RO Sol. Space

September 2009

"Mobility Header Home Agent Switch Message", [RFC 5142](#),
January 2008.

[Appendix A](#). Short Overview of all RO Solutions

This Appendix covers all four approaches (MNN to CN, MR to CN, MR to CR, MR to HA) to the RO problem with a short analysis.

[A.1](#). Analysis

For each solution class, a specific draft is shortly summarized before the analysis is performed.

[A.1.1](#). MNN to CN

In general, for this solution class information related to the MRs access network and access routers (ARs) is exposed to the MNNs. Based on this information, the MNNs can perform RO by themselves directly with the CN. The MR is therefore only in a supporting role.

[A.1.1.1](#). Proposal(s)

The MR is attached to a subnet with an appropriate AR. The draft [[I-D.ndproxy](#)] describes how the MR relays the subnet prefix of the AR inside the mobile network to the MNNs, that in turn use Autoconfiguration [[RFC4862](#)] to configure their IP addresses. The MR then acts as a ND proxy.

Another, similar approach relies on prefix delegation between MR and

AR [[I-D.pd](#)] where the AR receives a complete proper prefix that can be used inside the NEMO mobile network.

Packets sent over the R0 path use the Type 2 Routing Header and Home Destination Option as specified for R0 in [[RFC3775](#)].

[A.1.1.2](#). Analysis

This approach might face problems in the presence of Secure Neighbor Discovery in the access network [[RFC3971](#)] when using the MR as ND proxy. MNNs have to implement MIPv6 [[RFC3775](#)] for performing R0 themselves and the MR has to be upgraded as well.

Verifying the requirements against this solution approach, we come to the following conclusions:

1. Req1: Fulfilled. MNNs can decide by themselves to perform R0.
2. Req2: To be discussed. MNNs can configure one address per advertised (access network) prefix. The disadvantage is that the access network has to accept every MNN address as source address for packets, something that may not be supported if the

egress interface supports having only a single IP address.

3. Req3: Fulfilled. Can rely on standard MIPv6 R0; MR-HA tunnel therefore used for as long as R0 has not been completed.
4. Req4: Fulfilled. R0 is performed between the end systems - no additional single-point of failure for communication added.
5. Req5: Fulfilled. As long as the R0 path has not been established, packets can be sent over the MR-HA tunnel.
6. Req6: Problematic. Signaling overhead will be per MNN/CoA/home network prefix/CN. BGP not relevant.
7. Req7: Problematic. The signaling for R0 is equal to that of standard MIPv6. Several messages and RTTs are needed for every MNN that is performing R0.
8. Req8 #1: To be discussed. While the MNP itself is not part of

the R0 signaling, the addresses of the individual end systems within the R0 signaling is in cleartext. This is also the case for NEMO Basic support though, if no IPsec confidentiality protection is used for user data traffic.

9. Req8 #2: Problematic. If relying on standard MIPv6 R0, MNP/HoA verification can be broken.

10. Req9: To be discussed.

As conclusion, while this approach allows for high granularity of R0 triggering and setup due to the fact that the MNN is in charge, this approach has problems related to scalability and security if standard MIPv6 R0 is used at the MNNs.

[A.1.2.](#) MR to CN

In general, within this solution class the MR performs R0 on behalf of the MNN.

[A.1.2.1.](#) Proposal(s)

The MR acts as a transparent MIPv6-MN-proxy by performing standard MIPv6 R0 signaling on behalf of the MNN/LFN with the CN. The draft [[I-D.bernardos-nemo-miron](#)] describes the detailed operation where the MR performs the Return Routability procedure with its own CoA and the MNN address as HoA. Packets protected by IPsec AH between LFN and CN can not be supported in R0 mode, but instead have to be routed via the MR-HA tunnel.

Packets sent over the R0 path use the Type 2 Routing Header and Home Destination Option as specified for R0 in [[RFC3775](#)].

[A.1.2.2.](#) Analysis

The basic problem of this approach is the MR acting transparently between MNN and CN and performing the R0 as MN from the CN perspective. This can cause problems for security related protocols, as the MR actions can be regarded as man-in-the-middle attacks. This is particularly the case for IPsec AH.

Verifying the requirements against this solution approach, we come to

the following conclusions:

1. Req1: To be discussed. MR has to be configured with policies and has to perform packet inspection. Whether R0 can be specifically triggered for certain flows, depending on the traffic type, remains to be clarified though (esp. wrt IPsec).
2. Req2: Basically fulfilled. The MR could register several CoAs for the MNN with the help of [[I-D.draft-ietf-monami6-multiplecoa](#)], although no bulk registration is available in this case.
3. Req3: Fulfilled. Can rely on standard MIPv6 R0; MR-HA tunnel therefore used for as long as R0 has not been completed.
4. Req4: Fulfilled. R0 is performed with the correspondent node - no additional single-point of failure for communication added.
5. Req5: Fulfilled. As long as the R0 path has not been established, packets can be sent over the MR-HA tunnel.
6. Req6: Problematic. Signaling overhead will be per MNN/CoA/home network prefix/CN. BGP not relevant.
7. Req7: Problematic. The signaling for R0 is equal to that of standard MIPv6. Several messages and RTTs are needed for every MNN-CN pair for which R0 is performed.
8. Req8 #1: To be discussed. While the MNP itself is not part of the R0 signaling, the addresses of the individual end systems within the R0 signaling is in cleartext. This is also the case for NEMO Basic support though, if no IPsec confidentiality protection is used for user data traffic.
9. Req8 #2: Problematic. If relying on standard MIPv6 R0, MNP/HoA verification can be broken.

10. Req9: Problematic. IPsec AH not supported for reverse path from MNN/LFN to CN.

This approach allows to use simple LFNs as MNNs, but introduces problems due to the middlebox operation at the MR. Security concerns

with respect to the RO procedure itself are existing if standard MIPv6 RO is used. Scalability is similar to the approaches in [Appendix A.1.1](#).

[A.1.3](#). MR to CR

For this approach, a new mobility entity called Correspondent Router (CR) is introduced. The MR performs RO with the CR that forwards traffic from/to the CN/MNN.

[A.1.3.1](#). Proposal(s)

The CR approach was first introduced in [[I-D.orc](#)] as a possible solution to the NEMO RO problem. Later, a proposal based on the CR idea which specifically targeted the aeronautical requirements has been specified in [[I-D.draft-bernardos-mext-nemo-ro-cr](#)]. A third document [[I-D.draft-wakikawa-mext-cr-consideration](#)] has then been issued by the main author of [[I-D.orc](#)] for discussing possible issues and for providing considerations for the CR protocol.

The MR performs a discovery procedure to detect the CR that should be located either 1) on the routing-path to the CN or 2) within the network of the CN where it can announce proxy-routes via an IGP for the MNP(s) of the MR. Once the mobility binding has been established, the CR can intercept traffic (e.g. based on routes advertised via IGP) and tunnel it to the MR. Similarly, in the reverse direction the MR tunnels traffic destined to the prefix(es) served by the CR directly to this router. This is shown in Figure 1.

The discovery procedure is a critical part of the overall protocol and while the original document [[I-D.orc](#)] did propose a solution, issues related to this approach are discussed in [[I-D.draft-bernardos-mext-nemo-ro-cr](#)]. Hence, this topic has to be regarded as work in progress.

The authentication of the MNP in [[I-D.orc](#)] is achieved by adding a prefix sub-option containing the MNP(s) of the MR to the Home Test Init (HoTI) message. The HA only forwards the message to the CR if the originator of the message (MR) is also the owner of the prefix contained with the HoTI message.

CR Discovery is achieved by an ICMP-based mechanism similar to Dynamic Home Agent Address Discovery (DHAAD), based on the

Correspondent Nodes 64bit prefix.

The operation of a CR is more complicated if it is located in a multihomed site: asymmetric routing could result if the CR serving the MR on the forward path can not ensure to also intercept and forward packets to the MR on the reverse path.

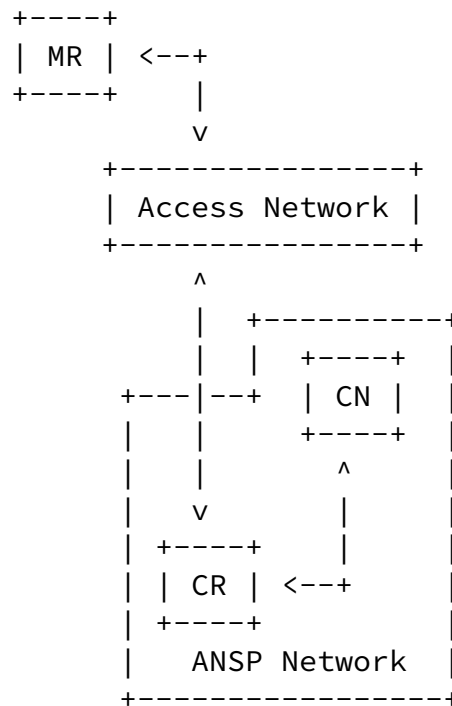


Figure 1: CR-based NEMO R0.

[A.1.3.2.](#) Analysis

The CR approach has both advantages and disadvantages.

1. Req1: To be discussed. MR has to be configured with policies and has to perform packet inspection. Whether R0 can be specifically triggered for certain flows, depending on the traffic type, remains to be clarified though (esp. wrt IPsec).
2. Req2: Basically fulfilled. The MR could register several CoAs for its MNP(s) with the help of [I-D.[draft-ietf-monami6-multiplecoa](#)], although no bulk registration is available in this case.
3. Req3: Fulfilled. Similar to standard MIPv6 R0; MR-HA tunnel therefore used for as long as R0 has not been completed.
4. Req4: To be discussed. The CR itself could be regarded as a new single point of failure, but if CR failure can be detected the

Internet-Draft

Aeronautical NEMO R0 Sol. Space

September 2009

MR-HA tunnel could be used again.

5. Req5: Fulfilled. As long as the R0 path has not been established, packets can be sent over the MR-HA tunnel. When the MR performs a transition to a new access network and an R0 state has been established with the CR, packets will be lost as the CR will send them to the old CoA of the MR. This is also true for NEMO Basic Support though, where the HA will send packets to the old CoA of the MR for as long as the mobility binding has not been updated.
6. Req6: To be discussed. Signaling overhead is per CoA/MNP/CR. BGP not relevant, IGP advertisements and routing table size at the CR grows with the number of registered MNPs though.
7. Req7: To be discussed. CNs within the same network can be served by a single CR, mobility signaling is therefore only performed for new CoAs. The overall signaling overhead is directly related to the distribution of CNs and CRs. An additional overhead is caused by the discovery procedure.
8. Req8 #1: To be discussed. The MNP itself is part of the R0 signaling and could be sent in cleartext.
9. Req8 #2: Problematic. If relying on standard MIPv6 R0, MNP/HoA verification can be broken.
10. Req9: Fulfilled. The tunneling between MR and CR preserves inner packet characteristics.

An advantage of CR is that R0 is scalable wrt the number of CNs as mobility signaling is performed per network or at least per grouped list of prefixes served by the CR (under the assumption that the CNs are located within those prefixes).

Another problem is the authentication between MR and CR: [[I-D.orc](#)] relies on the standard MIPv6 Return Routability procedure that has security weaknesses. In addition, the authentication of the CR to the MR is not taken into account in the original draft, but has to be considered as a potential threat.

[A.1.4.](#) MR to HA

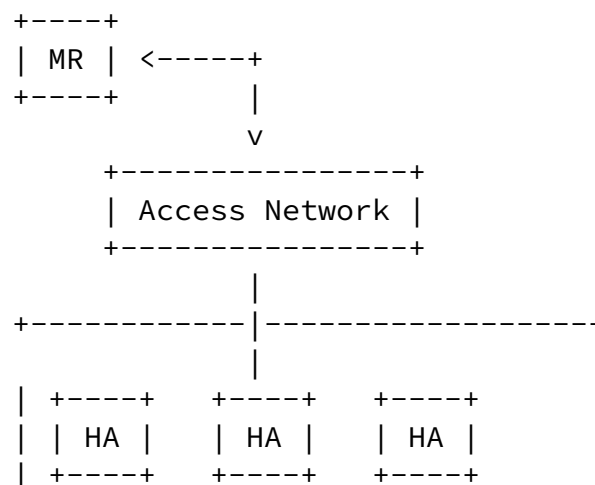
For this approach a new home network architecture is introduced where Home Agent functionality is shared among a set of instances that is geographically spread. At a given moment, based on the current topological location, the MR uses the closest HA instance.

[A.1.4.1.](#) Proposal

The Global Home Agent to Home Agent protocol [I-D.[draft-wakikawa-mext-global-haha-spec](#)] extends the original MIPv6/NEMO model, where only a single HA is available per MN, to an architecture where HAs are geographically distributed. MNs can bind to the closest HA to achieve a certain level of RO.

The Home Network relies on advertising a large common prefix via EGP that inflicts anycast routing. The traffic of a CN will be attracted by the topologically closest HA, just as mobility signaling from the MN will always be attracted by the topologically closest HA as well. In the latter case, the closer HA will inform the MNs current primary HA of the suboptimal routing. The primary HA will send a HA switch message that orders the MN to bind with the closer HA, based on signaling specified in [[RFC5142](#)]. This way, by reducing the distance between the MR and its HA, a certain degree of RO is achieved. Signaling between HAs is based on [I-D.[draft-ietf-mip6-hareliability](#)].

A graphical illustration is shown in Figure 2.



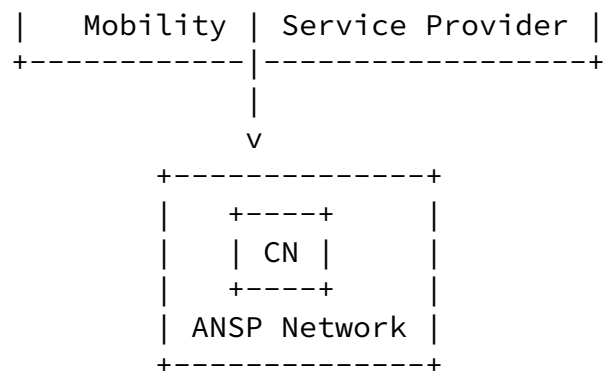


Figure 2: Global HA to HA.

[A.1.4.2.](#) Analysis

HAHA is significantly different from the other approaches as it does not require any mobility functionality in the CN or within the CN network. Mobility signaling, apart from the MR and its current primary HA, is performed within the various instances of the HAs in the ground network only.

1. Req1: To be discussed. RO is always implicitly provided by switching to a closer HA. Packet inspection at the MR would have to be introduced to identify the different flows. This requirement could only be discussed within the context of multihoming extensions, namely [\[I-D.draft-ietf-monami6-multiplecoa\]](#) and [\[I-D.draft-ietf-mext-flow-binding\]](#).
2. Req2: To be discussed. Given compatibility to [\[I-D.draft-ietf-monami6-multiplecoa\]](#) is provided, the MR could register several CoAs for its MNP(s). It is unclear though how the protocol could support simultaneous mobility signaling with two HAs if each one is considered to be close to each respective, active interface at the MR.
3. Req3: Fulfilled. As long as the MR has not registered to the new HA, the old MR-HA tunnel should be preserved.
4. Req4: Fulfilled. HAHA adds availability mechanisms to the home network as it is based on [\[I-D.draft-ietf-mip6-hareliability\]](#).

5. Req5: Fulfilled. As long as the binding with the new HA has not been completed, packets can be sent over the MR-previous HA tunnel.
6. Req6: To be discussed. Signaling overhead is per CoA/HA. HAHA relies on BGP advertisements to achieve anycast routing. The scale of the advertisements should only be per HA island/anycast prefix and not per MR/MNP.
7. Req7: Fulfilled. Mobility signaling is similar to MIPv6/NEMO and performed per MR and HA.
8. Req8 #1: Fulfilled. No exposure on the wireless link to what already happens for NEMO.
9. Req8 #2: Fulfilled. Security is on the same level as in NEMO.
10. Req9: Fulfilled. The tunneling between MR and HA preserves inner packet characteristics.

While the level of RO provided by HAHA is not as good as for the previous approaches, it can help eliminate continental round-trip times in the aviation scenario.

A large advantage are the strong security properties as mobility signaling is restricted to the MR and the HA (mobility service provider).

[A.2.](#) Applicability to the Aeronautical Environment

[A.2.1.](#) Overview

Table 1 provides a summary of the fulfillment of the individual requirements by each solution. Certain requirements turned out to be difficult to assess within the context of the solution - in that case the "D" categorization has been used to indicate that a more detailed investigation is needed on that subject.

As can be seen the first two approaches "MNN to CN" and "MR to CN" have problems related to scalability and efficient signaling, as RO signaling is always performed on a per CN basis. The latter, in

addition, has problems related to the security (address authentication) and adaptability requirements. The "MNN to CN" approach is mostly interesting from the nesting problem perspective only.

The rationale for Requirement 8 - Security (Section 3.8.1 in [\[I-D.ietf-mext-aero-reqs\]](#)) mentions that it is "reasonable to assume trust relationships between each MR and a number of mobility anchor points topologically near to its CNs". The rationale says that this trust relationship equates on having credentials for the authentication of the MNP between the mobility entities (MR, HA, CR). This statement actually rules out all "X to CN" approaches (under the assumption that standard MIPv6 Return Routability signaling should not be accepted due to its security limitations).

More promising solutions are the CR and the HAHA protocols and we have therefore focussed on these two approaches in more detail in [Section 3](#).

Requirement	MNN to CN	MR to CN	MR to CR	MR to HA
1	F	D	D	D
2	D	F	F	D
3	F	F	F	F
4	F	F	D	F
5	F	F	F	F
6	P	P	D	D

7	P	P	D	F
8-1	D	D	D	F
8-2	P	P	P	F
9	D	P	F	F

Abbreviations: F... Fulfilled, P... Problematic, D... To be discussed

Table 1: Overview of solution characteristics

Authors' Addresses

Christian Bauer
German Aerospace Center (DLR)

Email: Christian.Bauer@dlr.de

Serkan Ayaz
German Aerospace Center (DLR)

Email: Serkan.Ayaz@dlr.de

Arnauld Ebalard
EADS Innovation Works

Email: arnaud.ebalard@eads.net