

MEXT Working Group
Internet-Draft
Intended status: Experimental
Expires: January 8, 2009

C. Bernardos
M. Calderon
I. Soto
UC3M
July 7, 2008

Mobile IPv6 Route Optimisation for Network Mobility (MIRON)
draft-bernardos-mext-miron-00

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 January 8, 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 network may have. By extending the Mobile IPv6 support to Mobile Routers, nodes of the network are not required to support any kind of mobility, since packets must go through the Mobile Router-Home Agent (MRHA) bi-directional tunnel. Communications from/to a mobile network have to traverse the Home Agent, and therefore better paths may be available. Additionally, this solution adds packet overhead, due to the encapsulation.

This document describes an approach to the Route Optimisation for

NEMO, called Mobile IPv6 Route Optimisation for NEMO (MIRON). MIRON enables mobility-agnostic nodes within the mobile network to directly communicate (i.e. without traversing the MRHA bi-directional tunnel) with Correspondent Nodes. The solution is based on the Mobile Router performing the Mobile IPv6 Route Optimisation signalling on behalf of the nodes of the mobile network.

This document (in an appendix) also analyses how MIRON fits each of the currently identified NEMO Route Optimisation requirements for Operational Use in Aeronautics and Space Exploration.

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.	Protocol Overview	6
3.	Mobile Router operation	8
3.1.	Data Structures	8
3.2.	Performing Route Optimisation	9
4.	Home Agent, Local Fixed Node and Correspondent Node operation	10
5.	Conclusions	10
6.	Security Considerations	10
7.	IANA Considerations	11
8.	Acknowledgements	11
9.	References	11
9.1.	Normative References	11
9.2.	Informative References	12
Appendix A.	Analysis of MIRON and the Aeronautics requirements .	13
A.1.	Req1 - Separability	14
A.2.	Req2 - Multihoming	14
A.3.	Req3 - Latency	15
A.4.	Req4 - Availability	15
A.5.	Req5 - Packet Loss	16
A.6.	Req6 - Scalability	16
A.7.	Req7 - Efficient Signaling	17
A.8.	Req8 - Security	18
A.9.	Req9 - Adaptability	18
A.10.	Des1 - Configuration	19
A.11.	Des2 - Nesting	19
A.12.	Des3 - System Impact	19
A.13.	Des4 - VMN Support	20
A.14.	Des5 - Generality	20
Appendix B.	Change Log	20
	Authors' Addresses	20
	Intellectual Property and Copyright Statements	22

1. Introduction

This document assumes that the reader is familiar with the terminology related to Network Mobility [9] and [10] (Figure 1), and with the Mobile IPv6 [2] and NEMO Basic Support [3] protocols.

The goals of the Network Mobility (NEMO) Support are described in [11]. Basically, the main goal is to enable networks to move while preserving the communications of their nodes (Mobile Network Nodes, or MNNs), without requiring any mobility support on them. The NEMO Basic Support protocol [3] consists on setting a bi-directional tunnel (Figure 2) between the Mobile Router (MR) of the network (that connects the mobile network to the Internet) and its Home Agent (located at the Home Network of the mobile network). This is basically the Bi-directional Tunnel (BT) operation of Mobile IPv6 (MIPv6) [2], but without supporting Route Optimisation. Actually, the protocol extends the existing MIPv6 Binding Update (BU) message to inform the Home Agent (HA) of the current location of the mobile network (i.e. the MR's Care-of Address, CoA), through which the HA has to forward the packets addressed to the network prefix managed by the MR (Mobile Network Prefix, or MNP).

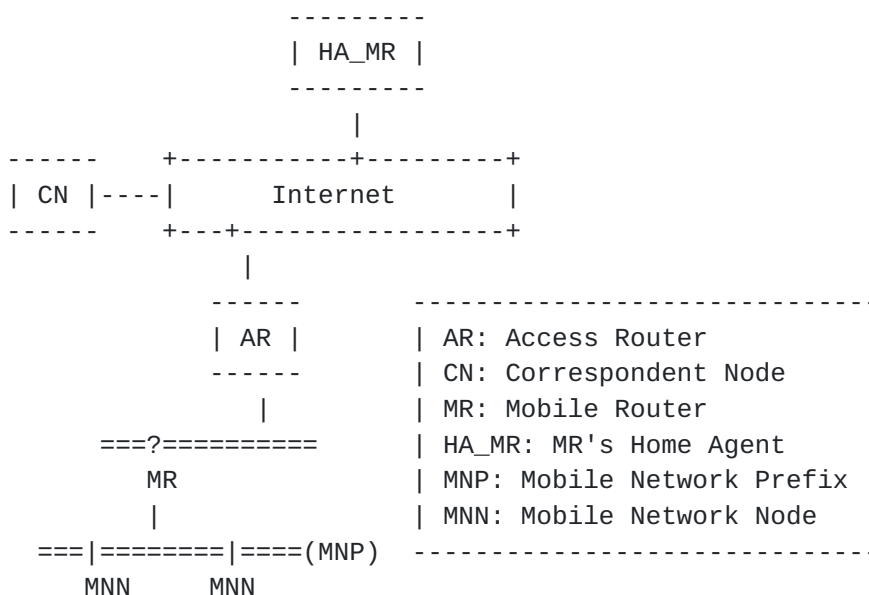


Figure 1: Basic scenario of Network Mobility

Because of the bi-directional tunnel that is established between HA and MR to transparently enable the movement of networks, the NEMO Basic Support protocol introduces the following limitations:

- o It forces suboptimal routing (known as angular or triangular routing), since packets are always forwarded through the HA following a suboptimal path and therefore adding a delay in the

packet delivery.

- o It introduces non-negligible packet overhead, reducing the Path MTU (PMTU). An additional IPv6 header (40 bytes) is added to every packet because of the MRHA bi-directional tunnel.
- o The HA becomes a bottleneck of the communication. This is because, even if a direct path is available between a MNN and a CN, the NEMO Basic Support protocol forces traffic to follow the CN-HA=MR-MNN path. This may cause the Home Link to be congested, resulting in some packets discarded.

In order to overcome such limitations, it is necessary to provide what have been called Route Optimisation for NEMO [12], [13], [14], [15]. In Mobile IPv6, the Route Optimisation is achieved by allowing the Mobile Node (MN) to send Binding Update messages also to the CNs. In this way the CN is also aware of the CoA where the MN's Home Address (HoA) is currently reachable. The Return Routability (RR) procedure is defined to authenticate new CoAs that the MN may use, thus securing the change that the CN makes in the IPv6 destination address (using the MN's CoA) of the packets it sends addressed to the MN's HoA [16].

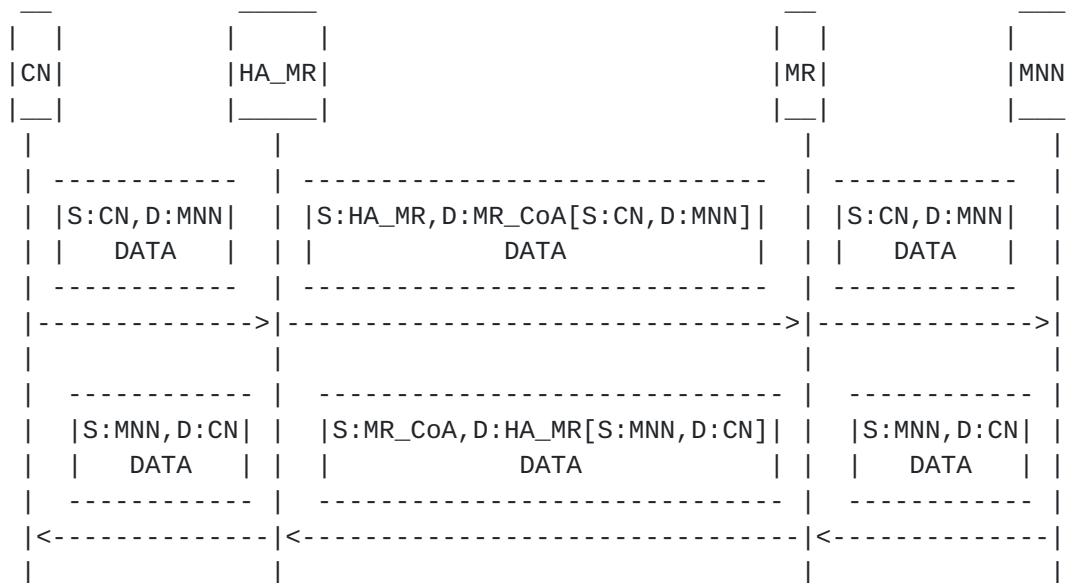


Figure 2: NEMO Basic Support protocol

This document describes a Route Optimisation solution for nodes of a mobile network that do not have (or use) any kind of mobility support, that is, the so-called Local Fixed Nodes (LFNs). The solution enables direct path communication between an LFN and a CN, without requiring any change on the operation of CNs nor LFNs. In order to do that, the MR performs all the Route Optimisation signalling and mobility tasks defined by Mobile IPv6 on behalf of the

LFNs attached to the mobile network.

2. Protocol Overview

The mechanism, called Mobile IPv6 Route Optimisation for NEMO (MIRON), essentially consists in enabling a MR to behave as a proxy for nodes that do not have any kind of mobility support (i.e. LFNs), performing the Mobile IPv6 Route Optimisation signalling and packet handling on behalf of the LFNs.

In order to enable packets to be directly routed between the CN and the LFN (avoiding the MRHA tunnel), the MR sends a MIPv6 Binding Update message on behalf of the LFN, binding the LFN's address to the MR's CoA.

Mobile IPv6 requires an additional security procedure to be performed before actually sending a Binding Update message to a certain CN (and therefore enabling the Route Optimisation between MN and CN). Basically, this procedure, called Return Routability (RR) -- needed in order to mitigate possible security concerns [16] -- is used to verify that the MN, besides being reachable at the HoA, is also able to send/respond to packets sent to a given address (different to its HoA). This mechanism can be deceived only if the routing infrastructure is compromised or if there is an attacker between the verifier node and the CoA (HoA and CoA) that are being verified. With these exceptions, the test is used to ensure that the MN's Home Address (HoA) and MN's Care-of Address (CoA) are collocated.

Since MIRON proposes the MR to behave as a "proxy" (Figure 3), the MR has to perform the Mobile IPv6 Return Routability procedure on behalf of the LFNs. This involves the MR sending the Home Test Init (HoTI) and Care-of Test Init (CoTI) messages to the CN and processing the replies (Home Test message HoT -- and Care-of Test message -- CoT). These messages are sent as specified in [2], using the LFN's address as the source address in the HoTI message -- which is sent encapsulated through the MR's HA --, and the MR's CoA as the source address in the CoTI message. With the information contained in the HoT and CoT messages, sent by the CN in response to the HoTI and CoTI messages respectively, the MR is able to build a BU message to be sent to the CN on behalf of the LFN. This message is sent using the MR's CoA as the packet source address and carries a Home Address destination option set to the LFN's address.

Once the Return Routability procedure has been done and the MR has sent the BU message -- binding the address of the LFN (belonging to the MR's MNP) to the MR's CoA -- packets between the CN and the LFN do not follow the suboptimal CN-HA=MR-LFN path anymore, but the CN-

MR-LFN optimised route (Figure 4). This signalling has to be repeated periodically, as specified by MIPv6 IPv6 (the RR procedure has to be repeated at least every 7 minutes, to avoid time-shifting attacks).

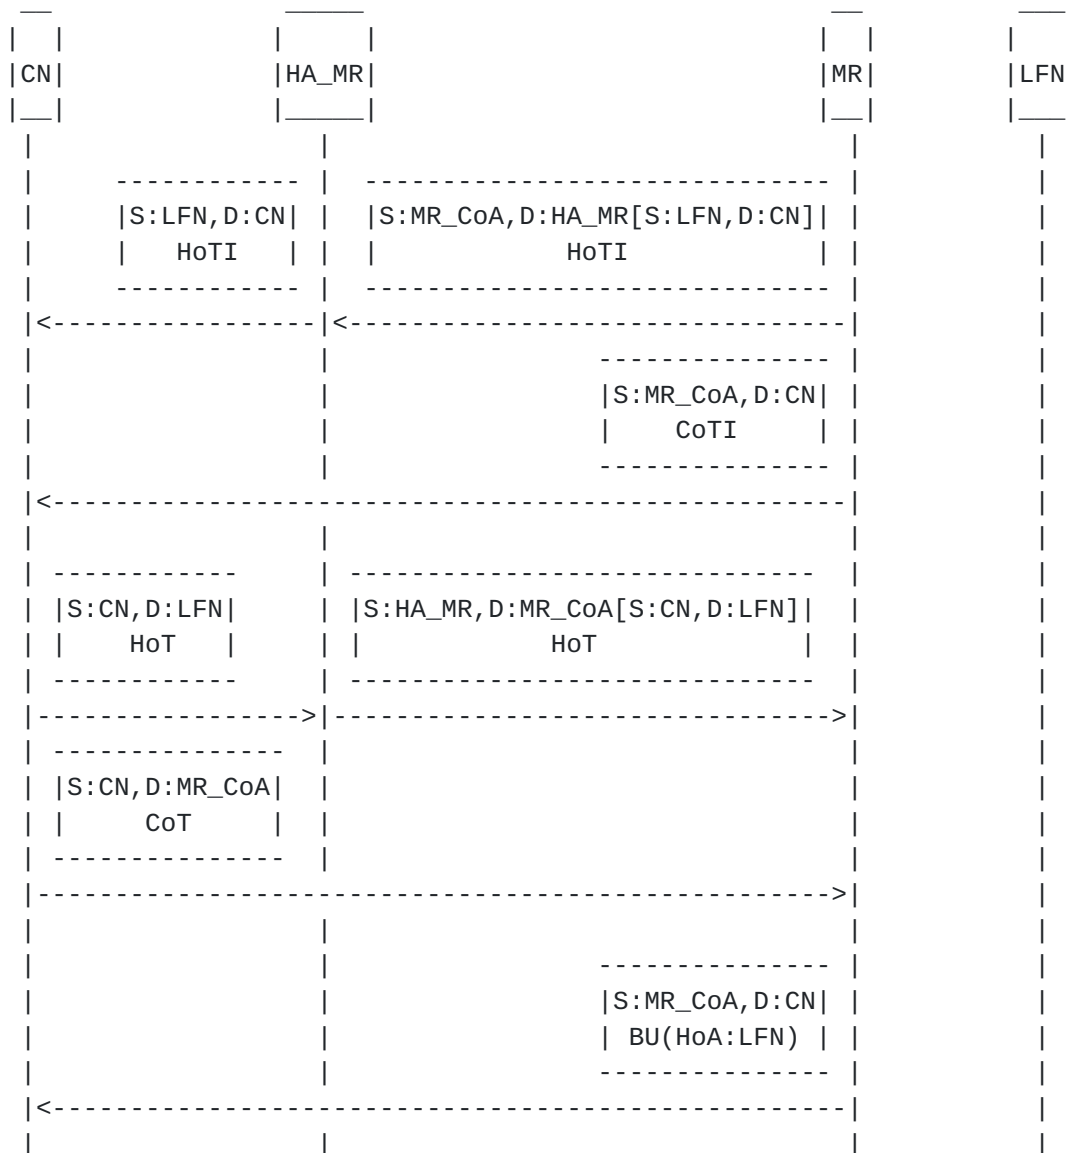


Figure 3: MIRON signalling

In addition to generating and receiving all the signalling related to Route Optimisation on behalf of the LFNs, the MR has also to process the "route optimised" packets sent by/directed to the LFNs (Figure 4):

- o Packets sent by a CN are addressed to the MR's CoA and contain IPv6 extension headers (a type 2 Routing Header) that are not understood by the LFN. Therefore, the MR has to change the

destination address and remove the routing header before delivering the packet to the LFN.

- o Packets sent by an LFN have to be also processed. The MR, in order to be able to send packets directly to the CN, has to use its CoA as source address of the packets and has also to add a Home Address destination option to every packet (set to the LFN's address).

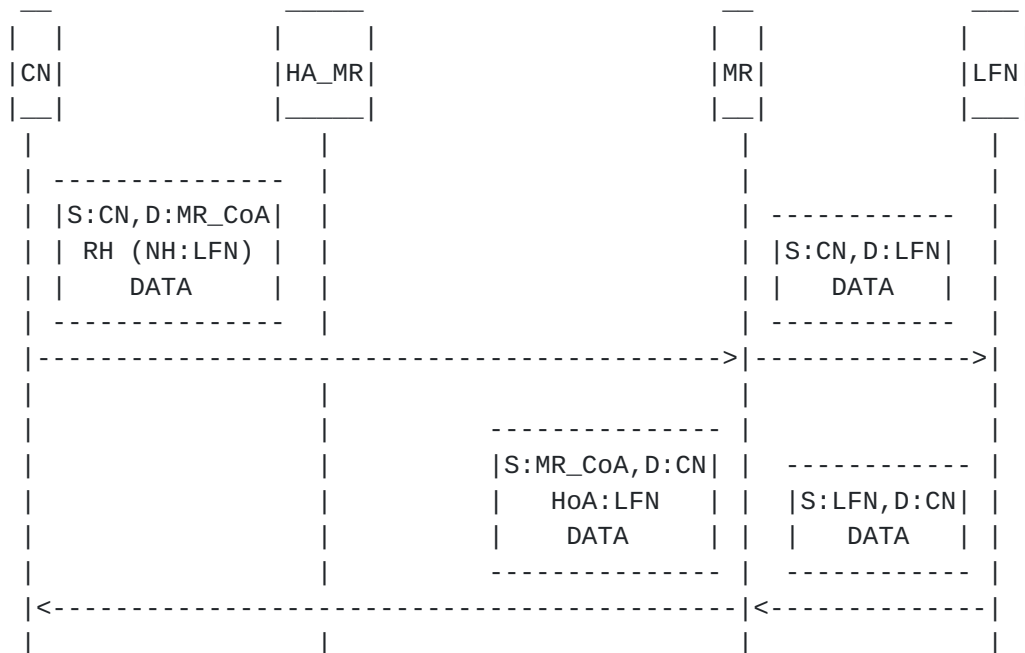


Figure 4: MIRON packet handling (Route Optimised operation)

3. Mobile Router operation

The Mobile Router operation defined by the NEMO Basic Support protocol [3] is extended in order to be able to generate and process the Mobile IPv6 Route Optimisation signalling (i.e. Return Routability and Binding Update) [2], since the MR is behaving as a "Proxy-MR" for their LFNs.

3.1. Data Structures

In addition to the data structures defined in [3], the MR need also to maintain the following information:

- o The MR extends the Binding Update List (BUL) to contain also some information per each LFN-CN communication that is being optimised. Basically the added fields are those described in [2] that are related to the Route Optimisation procedure defined by Mobile IPv6 (such as IP addresses of CNs, binding lifetimes, Return

Routability state, etc.), since the MR is behaving as Proxy-MR for the LFNs attached to it.

- o The BUL is not indexed only by the address of the CN, since there is a different binding per each CN-LFN pair. Therefore, the LFN's address has to be also included in every BUL entry.

3.2. Performing Route Optimisation

Since the proposed optimisation has to be done per each LFN-CN communication, the MR should track the different ongoing communications that attached LFNs may have, in order to identify potential LFN-CN communications that may be worth optimising. Due to the fact that optimising a certain LFN-CN communication involves a cost -- in terms of signalling and computation resources at the MR -- it may not be worth to perform such a optimisation to some kinds of flows (e.g., DNS queries). The decision about whether to perform Route Optimisation to a certain LFN-CN communication or not is out of the scope of this document.

Per each LFN-CN communication that has been decided to be route optimised, the MR has to perform the following actions:

- o The MR performs the Return Routability procedure as described in [Section 2](#).
- o The MR sends a MIPv6 Binding Update message to the CN on behalf of the LFN. The BU contains the address of the LFN as the MN's Home Address (HoA) and the MR's CoA as the MN's CoA (the MR's CoA is the only address that is reachable without requiring any agent to be deployed to forward packets to the current location of the mobile network). This procedure binds the LFN's address to the MR's CoA at the CN (i.e. an entry is added at the CN's Binding Cache).
- o The MR processes every packet received from the CN as follows:
 - * These packets carry the MR's CoA as destination address, and also carry a Type 2 Routing Header with the LFN's address as next hop. The MR processes the Routing Header of the packet, checking if the next hop address belongs to one of its LFNs and, if so, delivering the packet to the LFN.
- o The MR processes every packet received from the LFN as follows:
 - * The MR's CoA is set as IPv6 source address.
 - * An IPv6 Home Address destination option, carrying the address of the LFN, is inserted.

When MIRON is used to perform RO for a certain LFN-CN communication, the resulting PMTU of the path between LFN and CN might be bigger than the one when the MRHA tunnel is used. A MIRON MR implementation MAY decide to modify ICMPv6 Packet Too Big messages [\[4\]](#), [\[5\]](#) to include in the MTU field a value that takes into account the 24-byte overhead of MIRON (due to the Routing Header and Home Address

destination option) instead of the 40-byte overhead added when the MRHA bi-directional tunnel is used (due to the IPv6-in-IPv6 encapsulation).

4. Home Agent, Local Fixed Node and Correspondent Node operation

The operation of the Home Agent, the Local Fixed Nodes and the Correspondent Nodes remains unchanged. The only requirement is that CNs must implement the Correspondent Node part of the Route Optimisation defined by Mobile IPv6 (Section 9 of [2]).

5. Conclusions

This document describes a Route Optimisation for NEMO for LFN-CN flows. The MR performs all the Route Optimisation tasks on behalf of the LFNs that are attached to it. This involves generating and processing all the related signalling (that is, the Return Routability procedure and the Binding Update message), but also handling the packets sent and received by the LFNs, in order to enable the direct CN-MR-LFN route (avoiding the suboptimal CN-HA=MR-LFN path) without requiring any change on the operation of CNs nor LFNs.

The Route Optimisation here described is intended for LFN-CN communications in a non-nested NEMO. However, nothing prevents this solution to be also applied in a nested NEMO, avoiding the last tunnel, or even all the tunnels if the MR is provided somehow with a topologically valid IPv6 address that is reachable without traversing any HAs (for example, as described in [17]).

MIRON has been implemented [18] within the framework of the European DAIDALOS project (<http://www.ist-daidalos.org>). The implementation of the NEMO Basic Support protocol of MR and HA and the MIRON code at the MR have been done for Linux (2.6 kernel), mostly at user space (in C). The implementation used at the CNs is MIPL-2.0 (<http://www.mobile-ipv6.org>).

Conducted tests [17] show that the performance obtained with MIRON is much better -- in terms of latency and effective TCP throughput -- than the obtained by the NEMO Basic Support protocol, specially when the "distance" (i.e. RTT) between the HA and the CN is increased.

6. Security Considerations

Because an LFN trusts its MR for the routing of all its traffic (both

for inbound and outbound packets), allowing the MR to perform some signalling and processing on behalf of the LFNs attached to it does not introduce any new threat. From the architectural point of view, the solution is also natural, since the Route Optimisation support defined by Mobile IPv6 [2] conceptually could be implemented in multiple boxes. MIRON just applies this mechanism, by splitting the mobility functions among two different physical boxes, but actually the conceptual basis of the solution is the same as the one defined by Mobile IPv6.

7. IANA Considerations

This document has no actions for IANA.

8. Acknowledgements

Marcelo Bagnulo provided text for an earlier version of this draft.

The authors would like to thank Erik Nordmark and Pascal Thubert for their valuable comments. We also thank Antonio de la Oliva for implementing a first prototype of MIRON.

The work of Carlos J. Bernardos, Maria Calderon and Ignacio Soto described in this Internet-Draft is based on results of IST FP6 Integrated Projects DAIDALOS I & II. DAIDALOS I & II receive research funding from the European Community's Sixth Framework Programme. Apart from this, the European Commission has no responsibility for the content of this Internet-Draft. The information in this document is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.

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

9. References

9.1. 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.
 - [4] McCann, J., Deering, S., and J. Mogul, "Path MTU Discovery for IP version 6", [RFC 1981](#), August 1996.
 - [5] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", [RFC 4443](#), March 2006.
 - [6] Wakikawa, R., Devarapalli, V., Ernst, T., and K. Nagami, "Multiple Care-of Addresses Registration", [draft-ietf-monami6-multiplecoa-08](#) (work in progress), May 2008.
 - [7] Koodli, R., "Mobile IPv6 Fast Handovers", [RFC 5268](#), June 2008.
 - [8] Perkins, C., "Securing Mobile IPv6 Route Optimization Using a Static Shared Key", [RFC 4449](#), June 2006.

9.2. Informative References

- [9] Manner, J. and M. Kojo, "Mobility Related Terminology", [RFC 3753](#), June 2004.
- [10] Ernst, T. and H-Y. Lach, "Network Mobility Support Terminology", [RFC 4885](#), July 2007.
- [11] Ernst, T., "Network Mobility Support Goals and Requirements", [RFC 4886](#), July 2007.
- [12] Ng, C., Thubert, P., Watari, M., and F. Zhao, "Network Mobility Route Optimization Problem Statement", [RFC 4888](#), July 2007.
- [13] Ng, C., Zhao, F., Watari, M., and P. Thubert, "Network Mobility Route Optimization Solution Space Analysis", [RFC 4889](#), July 2007.
- [14] Zhao, F., "NEMO Route Optimization Problem Statement and Analysis", [draft-zhao-nemo-ro-ps-01](#) (work in progress), February 2005.
- [15] Thubert, P., Molteni, M., and C. Ng, "Taxonomy of Route Optimization models in the Nemo Context", [draft-thubert-nemo-ro-taxonomy-04](#) (work in progress), February 2005.

- [16] Nikander, P., Arkko, J., Aura, T., Montenegro, G., and E. Nordmark, "Mobile IP Version 6 Route Optimization Security Design Background", [RFC 4225](#), December 2005.
- [17] Calderon, M., Bernardos, C., Bagnulo, M., Soto, I., and A. de la Oliva, "MIRON: Mobile IPv6 Route Optimization for NEMO", IEEE Journal on Selected Areas in Communications (J-SAC), issue on Mobile Routers and Network Mobility, Volume 24, Number 9, pp. 1702-1716 , September 2006.
- [18] Bernardos, C., de la Oliva, A., Calderon, M., von Hugo, D., and H. Kahle, "NEMO: Network Mobility. Bringing ubiquity to the Internet access", IEEE INFOCOM 2006 demonstration, Barcelona , April 2006.
- [19] 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.
- [20] Ng, C., Ernst, T., Paik, E., and M. Bagnulo, "Analysis of Multihoming in Network Mobility Support", [RFC 4980](#), October 2007.
- [21] Dupont, F. and J. Combes, "Using IPsec between Mobile and Correspondent IPv6 Nodes", [draft-ietf-mip6-cn-ipsec-07](#) (work in progress), February 2008.
- [22] Baldessari, R., Ernst, T., and M. Lenardi, "Automotive Industry Requirements for NEMO Route Optimization", [draft-ietf-mext-nemo-ro-automotive-req-00](#) (work in progress), February 2008.
- [23] 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.

Appendix A. Analysis of MIRON and the Aeronautics requirements

This appendix looks at the Aeronautics requirements described in [19] and analyses how MIRON fits each of them. If a certain requirement cannot be fulfilled by MIRON as it is described in this document, possible modifications/extensions are also considered. This analysis aims at understanding if MIRON could be a good candidate to be used as a NEMO RO solution for the aeronautics scenario.

This appendix only analyses those requirements that involve LFNs, since the MIRON solution described here does only address the Route Optimisation of LFN-CN communications. For those requirements involving VMNs, an extended MIRON solution should be applied, such as the one described in [\[17\]](#).

[A.1.](#) Req1 - Separability

This requirement basically states that "an RO scheme MUST support configuration by a per-domain dynamic RO policy database. Entries in this database can be similar to those used in IPsec security policy databases in order to specify either bypassing or utilizing RO for specific flows".

This requirement is fulfilled by MIRON, since the Route Optimisation is performed in a per-flow basis and the decision about which flows are optimised is taken by the MR. Therefore, different approaches can be implemented in the MR (it is open to the particular MIRON implementation how to do it) to take this decision: static and dynamic policies (using a protocol to update MR's policies), decisions based on current load of the MR, etc.

[A.2.](#) Req2 - Multihoming

This requirement states that "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".

In MIRON, RO is achieved by the MR performing all the MIPv6-RO operations on behalf of connected LFNs. Therefore, MIRON can potentially benefit directly from any mechanism developed for MIPv6 to support multiple interfaces. For example, MIRON could use the Multiple-CoA mechanism [\[6\]](#) to enable an MR register different CoAs at CNs. It is also perfectly possible to support different MNPs for different domains, since the MR can manage the RO also with a per-MNP granularity.

We should also mention that although MIRON can benefit from multihoming solutions developed within the MEXT WG, multihoming issues in Network Mobility [\[20\]](#) should be tackled specifically by a general NEMO multihoming framework. Since MIRON does not modify in any way the NEMO Basic Support operation, it will also be compatible with such a general NEMO multihoming solution.

A.3. Req3 - Latency

This requirement says: "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".

This means that an RO solution MUST NOT prevent data packets from being forwarded through the MRHA bi-directional tunnel while the required RO operations are being performed. This requirement is also fulfilled by MIRON, since while the MR performs all the MIPv6-RO operations on behalf of LFNs, their communications still use the MRHA tunnel.

A.4. Req4 - Availability

This requirement states that "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". It is also stated that "an RO mechanism MUST NOT add any new single point of failure for communications in general".

This requirement requires special attention to be achieved by MIRON. On the one hand, current NEMO Basic Support protocol [3] does not fulfil that today, and therefore needs additional work to be carried-out. On the other hand, if we focus only on MIRON availability, the following are the potential points of failure that should be tackled:

- o Home Agent: a failure of the Home Agent -- in addition to disrupting communications that are being forwarded using the MRHA bi-directional tunnel -- could also cause Route Optimised communications to stop, since the Return Routability procedure (which is part of the MIPv6-RO mechanism) performed by the MR on behalf of LFNs requires the MRHA tunnel to be up. This Return Routability procedure should be repeated every seven minutes (MAX_RR_BINDING_LIFETIME=420 seconds) to avoid time-shifting attacks, and therefore, depending on how long a Home Agent is down and when the failure happens, route optimised communications might be affected. This issue is actually not introduced by MIRON, hence it could be avoided by using a general redundancy mechanism aimed for the NEMO Basic Support protocol.
- o Mobile Router: a failure of the MR would obviously disrupt established connections in a single-MR NEMO. In the case of a multiple-MR NEMO, additional mechanisms would be required in order to guarantee that route optimised communications managed by a particular MR would survive in case this MR fails. Although MIRON currently does not address this issue, additional mechanisms -- such as deploying back-to-back MRs in aircrafts or designing/reusing existing protocols to keep the RO state of several MRs synchronised -- might be used here.

- o Home Network reachability: in case of single-homed NEMOs, if the Home Network is not reachable, traffic through the MRHA bi-directional tunnel will be disrupted. Additionally, Return Routability checks performed by a MIRON MR on behalf of attached LFNs would also fail. Again, this issue is not introduced by MIRON, hence it needs to be addressed by a general redundancy mechanism suited for the NEMO Basic Support protocol.

A.5. Req5 - Packet Loss

This requirement says that "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".

It takes longer to finish a handover of a route optimised flow using MIRON than a normal NEMO Basic Support protocol handover, due to the additional RTT required to complete the MIPv6-RO signalling with CNs. If this additional RTT component in the overall handover delay is considered too much for a certain application, either this application could not use MIRON to provide NEMO RO or a make-before-break solution would be needed. Note that extending MIRON to support existing micromobility solutions -- such as Fast Handovers for Mobile IPv6 [7] -- would not require too much additional work, since MIRON basically re-uses standard MIPv6 mechanisms.

Alternatively, MIRON may decide to perform packet bi-casting and send route optimised traffic also through the MRHA bi-directional tunnel, during the time required to finish the MIPv6-RO signalling with CNs. That would allow not to lose any additional packets in the LFN-CN direction, but traffic from the CN would not be received by the LFN until the RO signalling has been completed. Therefore, this solution would only alleviate the problem for outbound packets and would require to take care of duplicated packets.

A.6. Req6 - Scalability

This requirement says that "an RO scheme MUST be simultaneously usable by the MNs 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".

Basically, there are three different aspects that may affect to the scalability of MIRON:

- o Memory consumption at the MR. MIRON needs some additional information to be stored at the MR, such extended BUL (since state information regarding each LFN-CN optimised pair is required). The required memory to store a BUL entry is relatively small and grows linearly with the number of LFN-CN route optimised pairs.
- o Processing load at the MR. MIRON requires the MR to perform some additional operations: inspection of every packet and special handling (that is, removal of the Routing Header in the CN-to-LFN direction and addition of the Home Address destination option in the LFN-to-CN direction) of route optimised packets. Regarding packet inspection, MIRON just needs to look at the source and destination addresses of every packet to track LFN-CN flows, so this inspection is quite similar to the normal inspection that a router does. Even if some local policies are implemented at the MR to enable smarter decisions about whether a certain flow should be optimised or not -- requiring the MR to look also at other fields in a packet (such as transport headers) -- this inspection is not much different than the inspection than typical firewall software does in a border (access) router.
- o Impact on the global routing system. MIRON does not have any impact on the global routing tables, and therefore it does not introduce any routing scalability issue, even with large deployments.

We can conclude that MIRON required resources grow linearly with the number of optimisations being performed, and that these required resources do not impose any constraint for modern available routers. Besides, MIRON does not impact in any way the global routing system.

A.7. Req7 - Efficient Signaling

This requirement is related to the additional signalling required by a NEMO RO solution, and basically states that "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".

With MIRON, in order to optimise a CN-LFN flow, the MR has to perform the MIPv6-RO signalling with the CN on behalf of the LFN. This signalling grows linearly with the number of CN-LFN pairs being optimised. In order to optimise an LFN-CN flow, the RR signalling (HoTI+CoTI+HoT+CoT) plus a Binding Update message should be sent every seven minutes. This means that 6.4 bps are required to keep each LFN-CN flow route optimised (without the MR moving). A NEMO that changes its point of attachment very frequently -- although this is not likely to happen in aircraft scenarios -- would require more signalling/support less optimised flows.

Another interesting metric is the data traffic overhead. MIRON introduces a 24-byte per packet overhead in every route optimised data packet, because of the Routing Header type 2 (added by the CN to the data packets sent to an LFN) and the Home Address destination option (added by the MR to the data packets sent by the LFN). Actually, the overhead introduced by MIRON is 16-byte less than the the overhead introduced by the NEMO Basic Support protocol, and therefore this means that MIRON reduces the overhead when compared with the NEMO non-optimised scenario.

A.8. Req8 - Security

This requirement is different depending on the considered traffic domain. For ATS/AOS domains, there are three sub-requirements: "a) The RO scheme MUST NOT further expose MNPs on the wireless link than already is the case for NEMO basic support; b) The RO scheme MUST permit the receiver of a BU to validate an MR's ownership of the CoAs claimed by an MR; and c) The RO scheme MUST ensure that only explicitly authorized MRs are able to perform a binding update for a specific MNP". For the PIES domain: "there are no additional requirements beyond those of normal Internet services and the same requirements for normal Mobile IPv6 RO apply".

MIRON does not meet -- without further modification, such as extending it to support solutions like [8], [21]-- requirements a) and c). In order to support those, RO signalling might need to be encrypted, thus requiring some security trust relationship between the MR and the CN (this is not considered as reasonable nowadays, though). On the other hand, MIRON supports requirement b), as this check is already performed by the RR procedure.

Regarding the requirements for the PIES domain, MIRON fulfils them, since this is already provided by MIPv6 RO security.

A.9. Req9 - Adaptability

This requirement states that "applications using new transport protocols, IPsec, or new IP options MUST be possible within an RO scheme".

In addition to RO decisions taken based on the lifetime of current communications (e.g., every data communication flow involving more packets than a particular threshold is route optimised), MIRON MAY make use of information about higher layer protocols to classify between flows that prefer the MRHA tunnel or a route optimised path. Depending on the mechanisms used to feed the decision intelligence mechanism at the MR, it might be required to perform some reconfiguration actions on MRs in order to enable new applications

and/or protocols benefit from RO. However, the use of unexpected/new higher layer protocols and/or applications would not make MIRON fail, but just revert on using the MRHA bi-directional tunnel for those flows that cannot be classified using existing filters/policies at the MR.

Regarding the particular issue of IPsec use, MIRON supports to route optimise communications that use IPsec ESP data traffic, since the ESP header comes after the Routing Header and Home Address destination option, and therefore MIRON does not affect ESP semantics. However, if IPsec AH is used in an LFN-CN communication, outbound (from the LFN to the CN) AH packets need to be reverse-tunnelled through the MRHA tunnel instead of being forwarded following the RO path, since adding the Home Address destination option would affect the AH integrity check. For inbound packets, there is no problem, since the MR just advances the source route and LFN can process these packets normally, without affecting AH semantics.

A.10. Des1 - Configuration

This requirement is not considered as a strict one, and basically states that "it is desirable that a NEMO RO solution be as simple to configure as possible and also easy to automatically disable if an undesirable state is reached".

MIRON configuration is not detailed in this document, since it is open to implementation. However, even if complex policies are used to determine which communication flows/applications are route optimised, MIRON configuration would be as simple as configuring today's firewalls. A MIRON MR does not require more configuration than a MIPv6 MN.

A.11. Des2 - Nesting

This requirement is not considered as a strict one, and basically states that "it is desirable if the RO mechanism supports RO for nested MRs".

MIRON, as it is described in this document, does not provide RO capabilities for nested MRs. However, it can be extended to support also these configurations. [17] describes one possible way of doing that.

A.12. Des3 - System Impact

This requirement is not considered as a strict one, and basically states that "low complexity in systems engineering and configuration

management is desirable in building and maintaining systems using the RO mechanism".

Since MIRON RO solution only requires changes on the MR, deploying MIRON is extremely easy in terms of global system impact. Only the MR is required to be configured, maintained and updated.

[A.13.](#) Des4 - VMN Support

This requirement is not considered as a strict one, and basically states that "it is strongly desirable for VMNs to be supported by the RO technique, but not strictly required".

As previously mentioned, this document has only described MIRON operation to optimise LFN-CN traffic flows. An extended MIRON version, such as the one described in [[17](#)], could be used to provide VMNs with Route Optimisation.

[A.14.](#) Des5 - Generality

This requirement is not considered as a strict one, and basically states that "an RO mechanism that is "general purpose", in that it is also readily usable in other contexts outside of aeronautics and space exploration, is desirable".

MIRON has been designed as a general NEMO RO framework, not being focused to address any particular scenario, but to be broadly-scoped. Therefore, MIRON could also be considered as a solution for other scenarios such as the vehicular [[22](#)] or consumer electronics ones [[23](#)].

[Appendix B.](#) Change Log

Changes from nemo-miron-01 to next-miron-00:

- o "Analysis of MIRON and the Aeronautics requirements" appendix updated as required by the latest available version of [[19](#)].
- o References updated.
- o Some text cleanup and minor changes.

Changes from nemo-miron-00 to nemo-miron-01:

- o Added appendix "Analysis of MIRON and the Aeronautics requirements".
- o Some text cleanup and minor 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

Maria Calderon
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain

Phone: +34 91624 8780
Email: maria@it.uc3m.es

Ignacio Soto
Universidad Carlos III de Madrid
Av. Universidad, 30
Leganes, Madrid 28911
Spain

Phone: +34 91624 5974
Email: isoto@it.uc3m.es

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.

