

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

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

Correspondent Router based Route Optimisation for NEMO (CRON)  
draft-bernardos-mext-nemo-ro-cr-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/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 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

Internet-Draft

CR-based RO for NEMO

July 2008

NEMO, based on the well-known concept of Correspondent Router. The solution aims at meeting the currently identified NEMO Route Optimisation requirements for Operational Use in Aeronautics and Space Exploration. Based on the ideas that have been proposed in the past, as well as some other extensions, this document describes a Correspondent Router based solution, trying to identify the most important open issues.

The main goal of this first version of the document is to describe an initial NEMO RO solution based on the deployment of Correspondent Routers and trigger the discussion within the MEXT WG about this kind of solution.

This document (in an appendix) also analyses how a Correspondent Router based solution 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].

Internet-Draft

CR-based R0 for NEMO

July 2008


## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction . . . . .</a>	<a href="#">4</a>
<a href="#">2.</a>	<a href="#">Protocol Overview . . . . .</a>	<a href="#">6</a>
<a href="#">3.</a>	<a href="#">Message Formats . . . . .</a>	<a href="#">7</a>
<a href="#">3.1.</a>	<a href="#">Correspondent Router Prefix Option . . . . .</a>	<a href="#">7</a>
<a href="#">4.</a>	<a href="#">Mobile Router operation . . . . .</a>	<a href="#">8</a>
<a href="#">4.1.</a>	<a href="#">Conceptual Data Structures . . . . .</a>	<a href="#">8</a>
<a href="#">4.2.</a>	<a href="#">Correspondent Router Discovery . . . . .</a>	<a href="#">9</a>
<a href="#">4.3.</a>	<a href="#">Correspondent Router Binding . . . . .</a>	<a href="#">10</a>
<a href="#">5.</a>	<a href="#">Correspondent Router operation . . . . .</a>	<a href="#">12</a>
<a href="#">5.1.</a>	<a href="#">Conceptual Data Structures . . . . .</a>	<a href="#">12</a>
<a href="#">5.2.</a>	<a href="#">Correspondent Router Binding . . . . .</a>	<a href="#">13</a>
<a href="#">5.3.</a>	<a href="#">Intercepting Packets from a Correspondent Node . . . . .</a>	<a href="#">13</a>
<a href="#">6.</a>	<a href="#">Security Considerations . . . . .</a>	<a href="#">14</a>
<a href="#">7.</a>	<a href="#">IANA Considerations . . . . .</a>	<a href="#">14</a>
<a href="#">8.</a>	<a href="#">Acknowledgements . . . . .</a>	<a href="#">14</a>
<a href="#">9.</a>	<a href="#">References . . . . .</a>	<a href="#">15</a>
<a href="#">9.1.</a>	<a href="#">Normative References . . . . .</a>	<a href="#">15</a>
<a href="#">9.2.</a>	<a href="#">Informative References . . . . .</a>	<a href="#">15</a>
<a href="#">Appendix A.</a>	<a href="#">Analysis of CRON and the Aeronautics requirements . . . . .</a>	<a href="#">16</a>
<a href="#">A.1.</a>	<a href="#">Req1 - Separability . . . . .</a>	<a href="#">16</a>
<a href="#">A.2.</a>	<a href="#">Req2 - Multihoming . . . . .</a>	<a href="#">17</a>
<a href="#">A.3.</a>	<a href="#">Req3 - Latency . . . . .</a>	<a href="#">17</a>
<a href="#">A.4.</a>	<a href="#">Req4 - Availability . . . . .</a>	<a href="#">18</a>
<a href="#">A.5.</a>	<a href="#">Req5 - Packet Loss . . . . .</a>	<a href="#">18</a>
<a href="#">A.6.</a>	<a href="#">Req6 - Scalability . . . . .</a>	<a href="#">19</a>
<a href="#">A.7.</a>	<a href="#">Req7 - Efficient Signaling . . . . .</a>	<a href="#">19</a>
<a href="#">A.8.</a>	<a href="#">Req8 - Security . . . . .</a>	<a href="#">20</a>
<a href="#">A.9.</a>	<a href="#">Req9 - Adaptability . . . . .</a>	<a href="#">20</a>
<a href="#">A.10.</a>	<a href="#">Des1 - Configuration . . . . .</a>	<a href="#">20</a>
<a href="#">A.11.</a>	<a href="#">Des2 - Nesting . . . . .</a>	<a href="#">21</a>
<a href="#">A.12.</a>	<a href="#">Des3 - System Impact . . . . .</a>	<a href="#">21</a>
<a href="#">A.13.</a>	<a href="#">Des4 - VMN Support . . . . .</a>	<a href="#">21</a>
<a href="#">A.14.</a>	<a href="#">Des5 - Generality . . . . .</a>	<a href="#">21</a>
	<a href="#">Authors' Addresses . . . . .</a>	<a href="#">22</a>

[1.](#) Introduction

This document assumes that the reader is familiar with the terminology related to Network Mobility [\[5\]](#) and [\[6\]](#) (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 [\[7\]](#). 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).



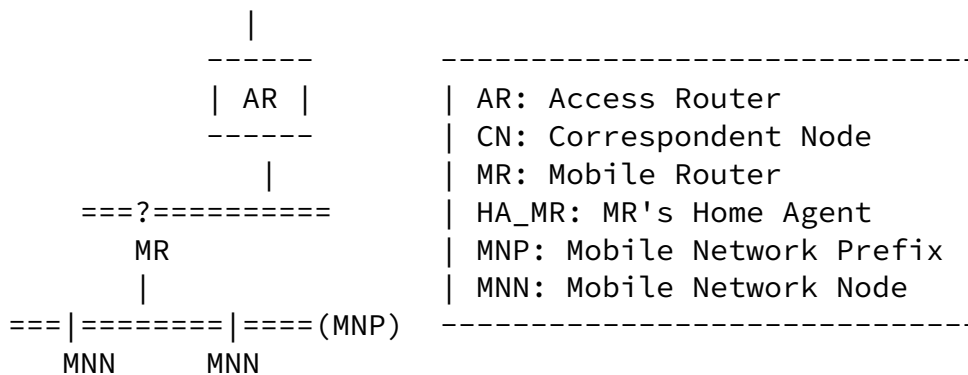


Figure 1: Basic scenario of Network Mobility


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 [8], [9], [10], [11].



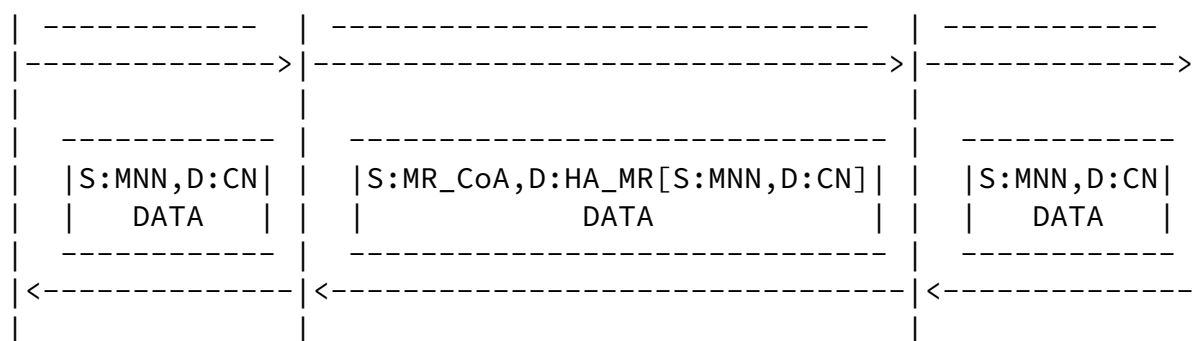


Figure 2: NEMO Basic Support protocol

This document collects some ideas about a Route Optimisation scheme based on optimising the routing path between an MNN and a CN by setting up a bi-directional tunnel between the MR of the NEMO and a router topologically close to the CN, called Correspondent Router (CR). Packets of the communication between the MNN and the CN could then use this bi-directional tunnel instead of the default MRHA one, in this way optimising the resulting routing path.

While this idea is far from being new ([12], [9]), the re-chartering of the MEXT WG (formerly NEMO WG) to specifically address and work on the design of a solution targeting a set of particular use case scenarios (namely aeronautics [13], vehicular [14] and consumer electronics [15]) makes worthwhile revisiting the idea of a CR-based NEMO RO solution and analyse whether this kind of approach would fit

the requirements of those scenarios. This document is a first attempt of doing so, focused on the requirements for Operational Use in Aeronautics and Space Exploration.

## 2. Protocol Overview

The approach described in this document, called Correspondent Router based Route Optimisation for NEMO (CRON), is based on the idea of extending the NEMO Basic Support protocol [3] to enable the MR managing bindings not only with the HA, but also with special routers, called Correspondent Routers (CRs), so the MRHA tunnel can be bypassed for traffic addressed to CNs that are topologically close to a CR.

CRON basically works as follows: an MR is forwarding packets to several communications between MNNs of the NEMO and CNs located in the Internet. For some of these communications, the MR may decide that it is better to try to avoid the use of the MRHA tunnel and try to set up a binding with a CR close to the other end-point of the communication (i.e. the CN). In order to do so, the MR needs to know whether there are any potential CRs that can be used and their IPv6 addresses. Once the MR knows that, it can select a CR among the potential ones (if any) and decide to establish a binding with it. To do so, a Binding Update/Binding Acknowledgement message exchange takes place. Once the signalling has occurred, the MR sets up a bi-directional tunnel with the CR and adds a routing entry towards the IPv6 prefix/address managed by the CR (the CN prefix), using the CR as next hop (through the established tunnel). Analogously, the CR sets up a routing entry towards the MNP/MNN IPv6 address, using the MR's CoA as next hop through the established tunnel. At this point, the traffic between the MNN (or set of MNNs with IPv6 addresses from a particular MNP) and the CN (or set of CNs with IPv6 addresses from a particular CN prefix) no longer traverses the HA, since the bi-directional tunnel established between the MR and the CR is used instead.

Next sections of this document describe in more detail the operation of CRON. It is not the goal of this first version to provide a complete and exhaustive specification (many details are not included), but rather to list the key operations that would be required for such a CR-based solution work, and, more importantly, identify the key requirements/assumptions/open issues that remain to be solved. This last point would be potentially very dependant on the considered NEMO RO use case.

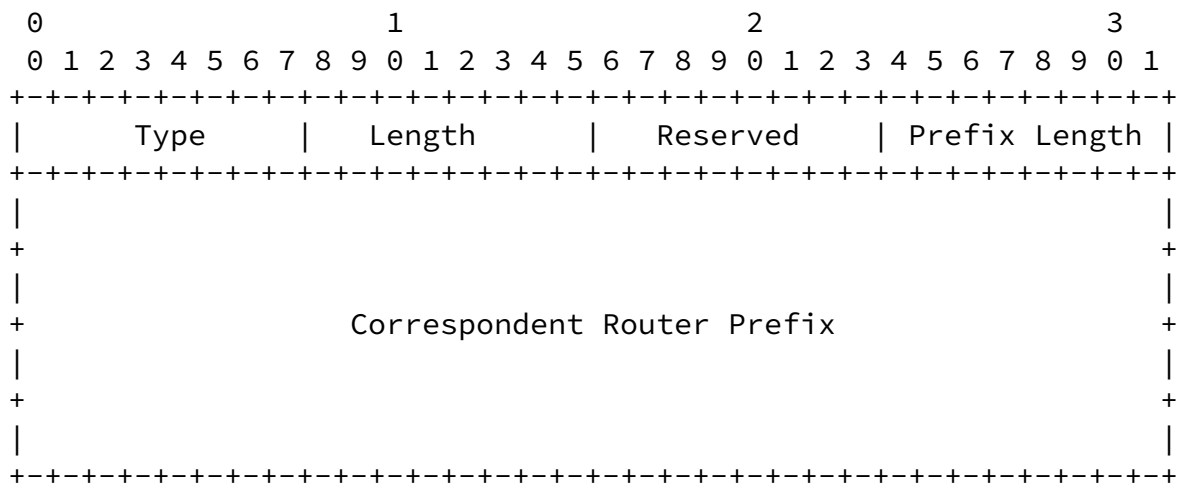
### [3.](#) Message Formats

#### [3.1.](#) Correspondent Router Prefix Option

The Correspondent Router Prefix Option is included in the Binding Updates to indicate to the Correspondent Router the CN prefix (could also be a host address) that is requested to be route optimised. This information and the prefix included in the Mobile Network Prefix

Option ([3]) is used by the CR to know which packets have to be sent through the MR-CR bi-directional tunnel. There could be multiple Correspondent Router Prefix Options if the CR is able to route optimise traffic from different CN prefixes and the MR wants traffic from more than one of these prefixes to be optimised. This option is also included by the CR in the BA messages.

The Correspondent Router Prefix Option has an alignment requirement of  $8n+4$ . Its format is as follows.



#### Type:

Number to be assigned by IANA.

#### Length:

Eight-bit unsigned integer indicating the length in octets of the option, excluding the type and length fields. Set to 18.

#### Reserved:

This field is unused for now. The value MUST be initialized to 0 by the sender and MUST be ignored by the receiver.

#### Prefix Length:

Eight-bit unsigned integer indicating the prefix length of the IPv6 prefix contained in the option.

#### Correspondent Router Prefix:



## 4. Mobile Router operation

The Mobile Router operation defined by the NEMO Basic Support protocol [3] is extended in order to be able to set up bindings with different CRs and establish bi-directional tunnels with them, used to route part of the traffic as indicated by the MR's policies.

The process of NEMO RO set-up MUST be initiated by the MR, and cannot be initiated by the CR. Next sections elaborate more on the different operations and mechanisms required to complete the whole NEMO RO.

### 4.1. Conceptual Data Structures

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

- o The MR extends the Binding Update List (BUL) to contain also some information per each communication that is being route optimised. Basically the added fields are the following:
  - \* A flag specifying whether the entry is a Correspondent Router entry or not.
  - \* Source IPv6 optimised prefix: this field contains the IPv6 address/prefix of those MNNs whose traffic is being optimised by using an MR-CR bi-directional tunnel with the CR. There could be multiple instances of this field in the BUL of the MR. These are basically the IPv6 addresses/prefixes included in the Mobile Network Prefix options carried in the BU/BA signalling exchanged by the MR and the CR.
  - \* Destination IPv6 optimised prefix: this field contains the IPv6 address/prefix of those CNs whose traffic is being optimised by using an MR-CR bi-directional tunnel with the CR. There could be multiple instances of this field in the BUL of the MR. These are basically the IPv6 addresses/prefixes included in the Correspondent Router Prefix options carried in the BU/BA signalling exchanged by the MR and the CR.
- o The MR SHOULD have a policy database that contains information regarding which communications should be route optimised and which not. It is up to the implementation the decision about the exact content of this database, as well as if it can be dynamically updated or it is pre-configured statically.

#### [4.2.](#) Correspondent Router Discovery

Prior to the binding establishment and tunnel set-up between an MR and a CR, the MR MUST find out whether there exist a CR that can be used to optimise the route between an attached MNN (or set of MNNs, having IPv6 addresses from the same MNP) and a CN (or set of CNs, having IPv6 addresses from the same IPv6 prefix). Besides, if there exist such a CR, the MR needs to know some information about it, such CR's IPv6 address.

This information may be known beforehand by the MR, by means of a static pre-configuration of the list of usable CRs and some additionally required information. Although it is up to the implementers to decide how this static information has to be processed and handled, the CR information MAY be indexed in a manner similar to a routing table, so when an MR has a target IPv6 address/prefix to which it wants to search for a candidate CR, it can perform a look-up in the CR database and get the most suitable CR (that is, the CR that is topologically closest to the target, using the longest-match prefix rule).

Although a static configuration might be enough for several scenarios, it seems desirable to support dynamic CR discovery. This is one of the hardest problems that a CR-based solution poses. There have been different proposals to tackle this problem. We next present the fundamentals of some of those, as well as some new ones:

- o If the CR is assumed to be always on the path between the MR to the CN (e.g., the CR is the gateway of the CN), the CR could inspect all received packets looking for a particular message. In this way, when there is a communication that is a candidate for being route optimised, the MR can send a BU message to the CN. If there is a CR available, the first CR on the path towards the CN would capture that packet (not forwarding it to the next hop) and process it accordingly. This approach seems to have severe limitations, since it assumes that available CRs will always be on the MR-CN path, but it deserves to be listed as it might be the case that this assumption is reasonable for some of the use cases currently considered by the MEXT WG.
- o Discovery based on/assisted by the DNS. The DNS service might be used in different ways, such as adding new type of registers, or including an 'A' register for the CR of each particular domain. In the latter case, the solution would work as follows: an MR attempting to optimise a certain MNN-CN communication would perform a reverse DNS query -- using the IPv6 address of the CN -- to obtain the qualified DNS name of the CN. Using the obtained

name, the MR would perform a query of the name 'CR.domainname', where 'domainname' is the domain part of the previously obtained

name. The 'A' register of 'CR.domainname' would contain (if that CR exists) the IPv6 address(es) of the CR(s) of that domain. This solution has also some drawbacks, since it assumes that each potential CN has a domain name published in the DNS, and assumes that all nodes belonging to a topological domain share the same domain names.

- o Use of a CR anycast address. The ORC draft [12] proposes the definition of anycast addresses for the CRs. In this way, the MR learns the CR anycast address from the CN's prefix and the defined anycast suffix. Using this anycast address, the MR sends CR discovery requests, waiting for CR discovery replies, in a way similar to the DHAAD mechanism of Mobile IPv6 [2]. This mechanism shares the limitations of DHAAD.
- o Deployment of CR-resolver servers. Given that some of the currently addressed use cases involve kind-of controlled environments, where certain trust relationships might be safely assumed, the deployment of a (potentially distributed, even in a hierarchical way) CR-resolver service may be considered. These CR-resolver servers would keep the a repository of CR-related information, including the IPv6 address of the CRs managing a particular IPv6 prefix/address. Every CR should keep that repository updated, so MRs can send queries to look for candidate CRs (the CR-resolver service would return the best CR, that is the one serving the longest-match prefix). In order to ensure the integrity of the stored information, CRs MUST have some type of security relationship with the CR-resolver service, so only authorised CRs can modify the stored information. Analogously, MRs SHOULD also have some trust relationship, so the MR is provided with some security guarantees regarding the authenticity of the information retrieved from CR-resolvers. It is also worth evaluating if it would be better that the HA is the one querying the CR-resolvers, and that the obtained information is provided to the MR through the HA (it might reduce the bandwidth consumption on the MR wireless access links and also reduce the number of required trust relationships). Both the requirement of deploying this CR-resolver service and the assumption of the existence of trust relationships should be carefully analysed for each of the currently addressed NEMO RO scenarios, in order to evaluate whether they are reasonable or not. This CR discovery approach

needs further work, TBD after gathering WG feedback about it.

#### 4.3. Correspondent Router Binding

Once the MR has learnt the IPv6 address of a CR that can be used to route optimise some traffic, it needs to perform the Binding Registration to this CR. The MR sends a Binding Update message protected with IPsec/IKE. It is assumed that MRs would have certificates that contain information regarding the Mobile Network

Bernardos, et al.

Expires January 8, 2009

[Page 10]

---

Internet-Draft

CR-based RO for NEMO

July 2008

Prefixes they are authorised to manage and their Home Address. It is also assumed that CRs would trust on those certificates, for example by means of a shared PKI infrastructure. Again, it needs to be discussed whether this assumption is too strong or it is reasonable for a given NEMO RO scenario (it seems that for AOS domain this requirement is less strong than for ATS domains [[16](#)]).

The Binding Update sent by the MR to the CR MUST have the Mobile Router (R) Flag set to 1, indicating to the CR that the BU is from an MR. The Home Registration (H) Flag MUST be set to 0. In this way, a CR that also implements the Home Agent functionality can differentiate Home Registration Binding requests from Correspondent Router Binding requests.

All BUs sent by MRs to CRs requesting a Correspondent Router Binding MUST carry at least a Mobile Network Prefix Option, containing the IPv6 prefix/address of the MNN(s) whose traffic is requested to be optimised. Therefore, only explicit binding mode is supported.

All BUs sent by MRs to CRs requesting a Correspondent Router Binding MUST carry at least a Correspondent Router Prefix Option, containing the IPv6 prefix/address of the CN(s) whose traffic is requested to be optimised. Traffic generated by nodes whose IPv6 addresses are not from the prefixes contained in these option, with destination the MNPs contained in the Mobile Network Prefix Option, MUST NOT be route optimised (the normal route MUST be used instead).

An MR MUST generate a different BU message per each MNN-CN (or MNP-CN Prefix) pair that wants to be optimised. Note that a binding is not restricted to the optimisation of traffic between an individual MNN and an individual CN, and can be established on an MNP-CN prefix basis. The CRON mechanism aims at providing the flexibility required

to support this granularity. Each BU contains the MNP prefix/MNN address information in the Mobile Network Prefix Options and the CN address/prefix information in the Correspondent Router Prefix Option.

The CR should be able to validate the CoA ownership claimed by the MR (that is, that the MR is actually reachable at its CoA). Depending on the considered use case, it might not be enough to rely on ingress filtering techniques at the access network. In those scenarios, the MR MUST test the CoA reachability following the Return Routability procedure (this would involve only the CoTI/CoT), as specified in [2].

Once the CR has received and processed the BU message, it SHOULD send a Binding Acknowledgement message to the MR. If no BA is received, the MR MAY retransmit BU messages as long as some rate limitation is applied. The MR MUST stop retransmitting when it receives a BA.

When an MR receives a BA message from a CR, the MR MUST validate the message according to some tests. This needs further elaboration (TBD in future versions of this document). Any BA not satisfying all of those tests MUST be silently ignored.

When an MR receives a packet carrying a valid BA, the MR MUST examine the Status field of the BA. If the status field indicates that the BU was accepted, the MR MUST update the corresponding entry in its Binding Update List to indicate that the Binding Update has been acknowledged. The MR MUST then set up a bi-directional tunnel with the IPv6 address of the CR as the other end-point of the tunnel, and add the entries to its routing table required to make use of the established tunnel in the forwarding of packets that match all of the following rules:

- o The IPv6 source address of the packet belongs to one of the IPv6 prefixes contained in the Mobile Network Prefix options included in the BA message. Note that this implies that source address based routing MAY be required in certain cases.
- o The IPv6 destination address of the packet belongs to one of the IPv6 prefixes contained in the Correspondent Router Prefix options included in the BA message.

## [5.](#) Correspondent Router operation

## [5.1.](#) Conceptual Data Structures

Correspondent Routers MUST maintain a Binding Cache (BC) of bindings with MRs. This BC is similar to the one kept by CNs with RO support as defined in [2]. In addition to the fields contained in the BC defined for a CN, the CR needs also to maintain the following information:

- o A flag specifying whether the entry is a Correspondent Router entry or not (applicable only on nodes which support also MIPv6 RO, to differentiate from normal MIPv6 BCEs).
- o Source IPv6 optimised prefix: this field contains the IPv6 address/prefix of those nodes managed by the CR (i.e. CNs) whose traffic is being optimised by using an MR-CR bi-directional tunnel with the MR. There could be multiple instances of this field in the BCE. These are basically the IPv6 addresses/prefixes included in the Correspondent Router Prefix options carried in the BU/BA signalling exchanged by the CR and the MR.
- o Destination IPv6 optimised prefix: this field contains the IPv6 address/prefix of those nodes (i.e. MNNs) whose traffic is being optimised by using an MR-CR bi-directional tunnel with the MR. There could be multiple instances of this field in the BCE. These are basically the IPv6 addresses/prefixes included in the Mobile

Network Prefix options carried in the BU/BA signalling exchanged by the CR and the MR.

- o The CR MAY have a policy database that contains information regarding which communications can be route optimised and which not. It is up to the implementation the decision about the exact content of this database, as well as if it can be dynamically updated or it is pre-configured statically.

## [5.2.](#) Correspondent Router Binding

When a CR receives a BU message from an MR, the CR MUST validate the message according to some tests. This needs further elaboration (TBD in future versions of this document). Any BU not satisfying all of those tests MUST be silently ignored. Additionally, some other test MUST be performed (such as authorisation check, etc.). If some of these tests fail, the CR SHOULD return a BA to the MR including an appropriate value in the Status field (TBD).

When a CR receives a packet carrying a valid BU, and all the tests are successful, the CR SHOULD add an entry on its BC and return a BA to the MR, setting the Status field to a value indicating success. This BA would also include all the Mobile Network Prefix and Correspondent Router Prefix Options included in the BU message. The CR MUST then set up a bi-directional tunnel with the MR's CoA as the other end-point of the tunnel, and add the entries in its routing table required to make use of this tunnel in the forwarding of packets that match all of the following rules:

- o The IPv6 source address of the packet belongs to the one of the IPv6 prefixes contained in the Correspondent Router Prefix options contained in the BA message. Note that this implies that source address based routing MAY be required in certain cases.
- o The IPv6 destination address of the packet belongs to the one of the IPv6 prefixes contained in the Mobile Network Prefix options included in the BA message.

### [5.3.](#) Intercepting Packets from a Correspondent Node

A Correspondent Router needs to intercept packets sent by all CNs from which it has a BCE, since the CR has to send those packets through the established MR-CR bi-directional tunnel. However, depending on the topological location of the CR, different operation may be needed in order to intercept packets from CNs.

If the CR is the gateway of all the CNs it manages -- that is, it is always on the path between all CNs and any other node outside the domain --, nothing needs to be done in order to intercept packets that have to be route optimised, since all packets are already received by the CR.

If the CR is not the gateway of all the CNs, in order to intercept those packets that need to be route optimised, the CR MUST inject routes advertising the corresponding MNPs within the domain of the CR. This would make routers in the domain to forward the packets addressed to those MNPs to the CR, so it can forward them to the right MRs using the established MR-CN bi-directional tunnel. Further attention is needed in order to come up with a flexible mechanism to enable a CR in a domain receive only those packets that need to be route optimised, without impacting on the routing of the rest of the packets -- that should follow the normal route (towards the respective home networks) --. Mechanisms enabling that behaviour are

subject of future versions of this document.

## 6. Security Considerations

CRON assumes that trust relationships exist between the MR/HA and the CR, allowing IPsec/IKE to be used to secure Binding Updates. Certificates are used to prove ownership of prefixes by MRs and CRs. The Return Routability mechanism is used by the MR to prove CoA ownership to the CR.

Given that CRON enables to change routing state at certain nodes, a more detailed security analysis is needed (TBD in a future version of this document).

## 7. IANA Considerations

This document requires IANA to assign a number for a new Mobility Option type (Correspondent Router Prefix).

## 8. Acknowledgements

This draft collects input from many previous works. The concept of Correspondent Router is known since at least 2002. Authors of the present document therefore would like to thank and acknowledge authors of these previous works.

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

Bernardos, et al.	Expires January 8, 2009	[Page 14]
-------------------	-------------------------	-----------

---

Internet-Draft	CR-based RO for NEMO	July 2008
----------------	----------------------	-----------

### 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] Wakikawa, R., Devarapalli, V., Ernst, T., and K. Nagami, "Multiple Care-of Addresses Registration", [draft-ietf-monami6-multiplecoa-08](#) (work in progress), May 2008.

## 9.2. Informative References

- [5] Manner, J. and M. Kojo, "Mobility Related Terminology", [RFC 3753](#), June 2004.
- [6] Ernst, T. and H-Y. Lach, "Network Mobility Support Terminology", [RFC 4885](#), July 2007.
- [7] Ernst, T., "Network Mobility Support Goals and Requirements", [RFC 4886](#), July 2007.
- [8] Ng, C., Thubert, P., Watari, M., and F. Zhao, "Network Mobility Route Optimization Problem Statement", [RFC 4888](#), July 2007.
- [9] Ng, C., Zhao, F., Watari, M., and P. Thubert, "Network Mobility Route Optimization Solution Space Analysis", [RFC 4889](#), July 2007.
- [10] Zhao, F., "NEMO Route Optimization Problem Statement and Analysis", [draft-zhao-nemo-ro-ps-01](#) (work in progress), February 2005.
- [11] 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.
- [12] Wakikawa, R., "Optimized Route Cache Protocol (ORC)", [draft-wakikawa-nemo-orc-01](#) (work in progress), November 2004.
- [13] 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.
- [14] 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.
- [15] 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.
- [16] Bauer, C. and S. Ayaz, "ATN Topology Considerations for Aeronautical NEMO R0", [draft-bauer-mext-aero-topology-00](#) (work in progress), July 2008.
- [17] Ng, C., Ernst, T., Paik, E., and M. Bagnulo, "Analysis of Multihoming in Network Mobility Support", [RFC 4980](#), October 2007.
- [18] Arkko, J. and I. Beijnum, "Failure Detection and Locator Pair Exploration Protocol for IPv6 Multihoming", [draft-ietf-shim6-failure-detection-13](#) (work in progress), June 2008.
- [19] Bernardos, C., Soto, I., Calderon, M., Boavida, F., and A. Azcorra, "VARON: Vehicular Ad-hoc Route Optimisation for NEMO0", Computer Communications, Volume 30, Issue 8, Pages 1765-1784 , June 2007.

## [Appendix A](#). Analysis of CRON and the Aeronautics requirements

This appendix looks at the Aeronautics requirements described in [\[13\]](#) and analyses how CRON fits each of them. If a certain requirement cannot be fulfilled by CRON as it is described in this document, possible modifications/extensions are also considered. This analysis aims at understanding if a CR-based solution could be a good candidate to be used as a NEMO R0 solution for the aeronautics scenario.

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

This requirement basically states that "an R0 scheme MUST support configuration by a per-domain dynamic R0 policy database. Entries in this database can be similar to those used in IPsec security policy

specific flows".

This requirement is fulfilled by CRON, since the Route Optimisation can be performed on an MNN-CN basis (MNP-CN prefix optimisations can also be performed) and the decision about which communications are optimised is taken by the MR. Therefore, different approaches can be implemented in the MR (it is open to the particular CRON 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 R0 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".

Since CRON achieves the NEMO R0 by performing bindings with different CRs, setting up bi-directional tunnels between a CR and an MR's CoA, it is possible to support multi-interfaced MRs and use different MNPs for different domains. Besides, CRON can potentially benefit directly from any mechanism developed for MIPv6 to support multiple interfaces (such as [\[4\]](#)).

We should also mention that although CRON can benefit from multihoming solutions developed within the MEXT WG, multihoming issues in Network Mobility [\[17\]](#) should be tackled specifically by a general NEMO multihoming framework. Since CRON 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 R0 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 R0 solution MUST NOT prevent data packets from being forwarded through the MRHA bi-directional tunnel while the

required R0 operations are being performed. This requirement is also fulfilled by CRON, since while the MR performs all the required signalling (i.e. CoTI/CoT and BU/BA), communications aimed at be route optimised still use the MRHA tunnel.

#### [A.4.](#) Req4 - Availability

This requirement states that "an R0 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 R0 mechanism MUST NOT add any new single point of failure for communications in general".

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, CRON brings the role of the Correspondent Router into the picture. Therefore, a new potential point of failure is added: the CR.

If a CR fails, communications being optimised would obviously be disrupted. Although CRON currently does not address this issue, additional mechanisms -- such as deploying several redundant or back-to-back CRs, or designing/reusing existing protocols to keep the R0 state of several CRs synchronised -- might be used here. In any case, by using short lifetimes, the potential negative effect of such a CR failure may be reduced. Besides, protocols such as REAP [\[18\]](#) can be used to effectively detect failures between the MR and the CR (not only caused by a CR failure, but also by a problem on the path between them).

In case an MR fails, if it is the only one deployed at the NEMO, this would obviously disrupt established connections. 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. Again, although CRON 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 R0 state of several MRs synchronised

-- might be used here.

#### [A.5.](#) Req5 - Packet Loss

This requirement says that "an R0 scheme SHOULD NOT cause either loss or duplication of data packets during R0 path establishment, usage, or transition, above that caused in the NEMO basic support case. An R0 scheme MUST NOT itself create non-transient losses and duplications within a packet stream".

The use of CRON does not add any significant delay nor causes any additional packet loss compared to the normal NEMO Basic Support protocol handover. It is worth to mention that some signalling is required to update the bindings at the CRs, and although this can be done in parallel with the Home Registration, a small delay can be

introduced because of this. Micro-mobility techniques can be used if required to mitigate that effect, if required.

#### [A.6.](#) Req6 - Scalability

This requirement says that "an R0 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 R0".

There are different aspects that may affect to the scalability of CRON:

- o Memory consumption at the MR. CRON needs some additional information to be stored at the MR, such extended BUL. The required memory to store this extended BUL is relatively small and grows linearly with the number of different route optimised communications.
- o Processing load at the MR. CRON requires more complex routing operations at the MR. The routing table of an MR performing R0 operations would have more entries than a normal [RFC 3963](#) MR, the MR should set-up and maintain more bi-directional tunnels, and source address based routing might be required. However, all of these operations do not add significant load, and in any case complexity grows linearly with the number of different route

- optimised communications.
- o Processing load and memory consumption at the CR. The same reasoning applies here. Besides, multiple CRs can be deployed on sites supporting a high load of route optimised traffic.
  - o Impact on the global routing system. CRON does not have any impact on the global routing tables, and therefore it does not introduce any routing scalability issue, even with large deployments.

#### [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 CRON, the amount of required signalling depends on the number and type of communications that are optimised. Since almost any granularity is supported by CRON, in those cases in which individual MNN-CN optimisations are not required -- and MNP-CN prefix ones are used instead -- the amount of signalling needed is very small. In

order to perform an optimisation (generally speaking, a MNP-CN prefix optimisation), a BU/BA message exchange is required (plus probably a CoTI/CoT exchange to test MR's CoA reachability).

#### [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".

CRON meets all aforementioned requirements. a) Since BU/BA signalling is sent protected with IPsec, MNPs are not exposed, b) the MR's CoA

ownership can be tested by means of the RR procedure (in particular, by means of the CoTI/CoT exchange), and c) by the use of certificates and local policies, only authorised MRs can perform a binding for a specific MNP.

#### [A.9.](#) Req9 - Adaptability

This requirement states that "applications using new transport protocols, IPsec, or new IP options MUST be possible within an R0 scheme".

Since CRON makes use of a bi-directional tunnel between the MR and the CR, encapsulating data packets into the tunnel -- without performing any modifications to them --, this requirement is also met.

#### [A.10.](#) Des1 - Configuration

This requirement is not considered as a strict one, and basically states that "it is desirable that a NEMO R0 solution be as simple to configure as possible and also easy to automatically disable if an undesirable state is reached".

CRON configuration is not detailed in this document, since it is open to implementation. However, even if complex policies are used to determine which communication are route optimised, CRON configuration would be as simple as configuring today's firewalls. Some coordination is needed though to configure properly MRs, CRs -- and likely additional infrastructure (such as CR-resolvers) -- in order

to effectively enable R0 to take place.

#### [A.11.](#) Des2 - Nesting

This requirement is not considered as a strict one, and basically states that "it is desirable if the R0 mechanism supports R0 for nested MRs".

CRON, as it is described in this document, does not provide R0 capabilities for nested MRs.

#### [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 R0 mechanism".

CRON requires changes on the MR, and the deployment of additional entities: the CRs (although this might be done by simply updating the software of some existing routers at the CN domains to support the CR functionality). The deployment of CRON also requires trust relationships between MRs/HAs and CRs to be in place. In order to help MRs in the discovery of CRs, the deployment of external services, such as CR-resolvers, might be required.

#### [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 R0 technique, but not strictly required".

CRON is aimed at bypassing the MR's HA, but it does not avoid the traversal of VMNs' HAs. Therefore, a completely optimal path is not follow by packets of VMNs (only the MR's HA is bypassed).

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

This requirement is not considered as a strict one, and basically states that "an R0 mechanism that is "general purpose", in that it is also readily usable in other contexts outside of aeronautics and space exploration, is desirable".

Correspondent Router approaches were designed as general NEMO R0 frameworks, not being focused to address any particular scenario. Current CRON design has been tailored to address the particular scenario of aeronautics and space exploration. However, CRON -- with or without modifications/extensions -- could also be considered as a

solution for other scenarios such as the vehicular [\[14\]](#) one. Besides, extensions to enable the use of CRON without the requiring the deployment of certificates -- using techniques similar to the ones described in [\[19\]](#) -- are currently being analysed and will be included in future revisions of this document. These extensions



would enable extending the applicability of CRON to other scenarios such as the consumer electronics one [[15](#)].

#### 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)

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

Phone: +34 91624 8780  
Email: [maria@it.uc3m.es](mailto: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](mailto: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](mailto:ietf-ipr@ietf.org).

