

MEXT Working Group
Internet-Draft
Intended status: Informational
Expires: February 22, 2008

W. Eddy
Verizon
W. Ivancic
NASA
T. Davis
Boeing
August 21, 2007

NEMO Route Optimization Requirements for Operational Use in Aeronautics
and Space Exploration Mobile Networks
[draft-eddy-nemo-aero-reqs-02](#)

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 February 22, 2008.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This document describes the requirements and desired properties of NEMO Route Optimization techniques for use in global networked communications systems for aeronautics and space exploration.

This version has been review by members of the International Civil Aviation Organization (ICAO) and other aeronautical communications standards bodies.

Table of Contents

1.	Introduction	3
2.	NEMO RO Scenarios	5
2.1.	Aeronautical Communications Scenarios	5
2.2.	Space Exploration Scenarios	9
3.	Required Characteristics	11
3.1.	Req1 - Separability	11
3.2.	Req2 - Multihoming	12
3.3.	Req3 - Latency	13
3.4.	Req4 - Availability	13
3.5.	Req5 - Packet Loss	14
3.6.	Req6 - Scalability	15
3.7.	Req7 - Efficient Signaling	15
3.8.	Req8 - Security	16
3.9.	Req9 - Adaptability	17
4.	Desirable Characteristics	17
4.1.	Des1 - Configuration	17
4.2.	Des2 - Nesting	18
4.3.	Des3 - System Impact	18
4.4.	Des4 - VMN Support	18
4.5.	Des5 - Generality	19
5.	Security Considerations	19
6.	IANA Considerations	19
7.	Acknowledgments	19
8.	Changes from draft-eddy-nemo-aero-reqs-01	20
9.	Informative References	21
Appendix A.	Basics of IP-based Aeronautical Networking	22
Appendix B.	Basics of IP-based Space Networking	23
	Authors' Addresses	23
	Intellectual Property and Copyright Statements	25

1. Introduction

As background the NEMO terminology and NEMO goals and requirements documents are suggested reading [1] [2]. The drafts produced as part of the Mobile Platform Internet (MPI) effort are also of relevance, and some of the material in this document is borrowed from the MPI drafts [3] [4].

The base NEMO standard [5] allows Mobile IPv6 [6] to be used by mobile routers, although NEMO does not support Mobile IPv6's Route Optimization (RO) features. NEMO allows mobile networks to efficiently remain reachable via fixed IP address prefixes no matter where they relocate within the network topology. This is accomplished through the maintenance of a bidirectional tunnel between a NEMO Mobile Router (MR) and a NEMO Home Agent (HA) placed at some relatively stable point in the network. Corresponding Nodes (CNs) that communicate with Mobile Network Nodes (MNNs) behind an MR do so through the HA and MRHA tunnel in both directions. Since the use of this tunnel may have significant drawbacks [7], RO techniques that allow a more direct path between the CN and MR to be used are highly desirable.

For decades, mobile networks of some form have been used for communications with people and avionics equipment onboard aircraft and spacecraft. These have not typically used IP, although architectures are being devised and deployed based on IP in both the aeronautics and space exploration communities (see [Appendix A](#) and [Appendix B](#) for more information). An aircraft or spacecraft generally contains many computing nodes, sensors, and other devices that are possible to address individually with IPv6. This is desirable to support network-centric operations concepts. Given that a craft has only a small number of access links, it is very natural to use NEMO MRs to localize the functions needed to manage the large onboard network's reachability over the few dynamic access links. On an aircraft, the nodes are arranged in multiple independent networks, based on their functions. These multiple networks are required for regulatory reasons to have different treatments of their air-ground traffic, and often must use distinct air-ground links and service providers.

Many properties of the aeronautics and space environments amplify the known issues with NEMO bidirectional MRHA tunnels [7] even further.

Longer routes leading to increased delay and additional infrastructure load - In aeronautical networks (e.g. using Plain Old Aircraft Communication Addressing and Reporting System (ACARS) or ACARS over VHF Data Link (VDL) mode 2) the queueing delays are often long due to ARQ mechanisms and underprovisioned radio links.

Furthermore, for aeronautical communications systems that pass through geosynchronous satellites, and for space exploration, the propagation delays are also long. These delays combined with the additional need to cross continents in order to transport packets between ground stations and CNS means that delays are already quite high in aeronautical and space networks without the addition of an MRHA tunnel. The increased delays caused by MRHA tunnels may be unacceptable in meeting Required Communication Performance [8].

Increased packet overhead - Given the constrained link bandwidths available in even future communications systems for aeronautics and space exploration, planners are extremely adverse to header overhead. Since any amount of available link capacity can be utilized for extra situational awareness, science data, etc., every byte of header/tunnel overhead displaces a byte of useful data.

Increased chances of packet fragmentation - Packet fragmentation exacerbates the issue with header overhead and adds to unreliability and possibly the need for end-to-end retransmission if fragments are lost. Since error-prone radio links are sometimes used in the aeronautical and space environments, loss of fragments resulting in loss of entire packets is possible. The odds of packets requiring fragmentation must be minimized in aeronautical and space networks.

Increased susceptibility to failure - The additional likelihood of either a single link failure disrupting all communications or an HA failure disrupting all communications is problematic when using MRHA tunnels for command and control applications that require high availability for safety-of-life or safety-of-mission.

For these reasons, an RO extension to NEMO is highly desirable for use in aeronautical and space networking. In fact, a standard RO mechanism may even be necessary before some planners will seriously consider advancing use of the NEMO technology from experimental demonstrations to operational use within their communications architectures. Without an RO solution, NEMO is difficult to justify for realistic operational consideration.

In [Section 2](#) we describe the relevant high-level features of the access and onboard networks envisioned for use in aeronautics and space exploration, as they influence the properties of usable NEMO RO solutions. [Section 3](#) then lists the technical and functional characteristics that are absolutely required of a NEMO RO solution for these environments, while [Section 4](#) lists some additional characteristics that are desired, but not necessarily required. In

[Appendix A](#) and [Appendix B](#) we provide brief primers on the specific operational concepts used in aeronautics and space exploration respectively for IP-based network architectures.

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](#). Although this document does not specify an actual protocol, but rather just the requirements for a protocol, it still uses the [RFC 2119](#) language to make the requirements clear.

[2.](#) NEMO RO Scenarios

To motivate and drive the development of the requirements and desirable features for NEMO RO solutions, in this section some operational characteristics are described to explain how access networks, HAS, and CNs are configured and distributed geographically and topologically in aeronautical and space network architectures.

[2.1.](#) Aeronautical Communications Scenarios

Since aircraft may be simultaneously connected to multiple ground access networks using diverse technologies with different coverage properties, it is difficult to say much in general about the rate of changes in active access links and care-of addresses (CoAs). As one data point, for using VDL mode 2 data links, the length of time spent on a single access channel varies depending on the stage of flight. On the airport surface, VDL mode 2 access is stable while a plane is unloaded, loaded, refueled, etc., but other wired and wireless LAN links (e.g. the Gatelink system) may come and go. Immediately after takeoff and before landing, planes are in the terminal maneuvering area for approximately 10 minutes and stably use another VDL mode 2 channel. During en-route flight, handovers between VDL mode 2 channels may occur every 30 to 60 minutes, depending on the exact flight plan and layout of towers, cells, and sectors used by a service provider.

The characteristics of a data flow between a CN and MNN varies both depending on the data flow's domain, and the particular application within the domain. Even within the three aeronautical domains described below, there are varying classes of service that are regulated differently, but this level of detail has been abstracted out for the purposes of this document. It is assumed that any viable NEMO RO solution will be able to support a granularity of configuration with many sub-classes of traffic within each of the specific domains listed here.

2.1.1.1. Air Traffic Services Domain

The MNNs involved in Air Traffic Services (ATS) consist of pieces of avionics hardware onboard an aircraft used to provide navigation, control, and situational awareness. The applications run by these MNNs are mostly critical to the safety of the lives of the passengers and crew. The MNN equipment may consist of a range of devices from typical laptop computers to very specialized avionics devices. These MNNs will mostly be Local Fixed Nodes (LFNs), with a few Local Mobile Nodes (LMNs) to support Electronic Flight Bags, for instance. It can be assumed that Visiting Mobile Nodes (VMNs) are never used within the ATS domain.

An MR used for ATS will be capable of using multiple data links (at least VHF-based, satellite, HF-based, and wired), and will likely be supported by a backup unit in the case of failure, leading to a case of a multi-homed MR that is at least multi-interfaced and possibly multi-prefixed as well, in NEMO terminology.

Since for ATS, the MRs and MNNs are under regulatory control and are actively tested and maintained, it is not unreasonable to assume that special patches or software be running on these devices to enable NEMO RO. In fact, since these devices are accessed by skilled technicians and professionals, it may be acceptable even if complex configuration is required for NEMO RO. Of course simplicity in setup and configuration is desirable, however.

Data flows from the ATS domain may be assumed to consist mainly of short transactional exchanges such as clearance requests and grants. Future ATS communications are likely to include longer messages and higher message frequencies for positional awareness and trajectory intent of all vehicles in motion at the airport and all aircraft within a thirty mile range during flight. Many of these may be aircraft-to-aircraft, but the majority of current exchanges are between the MNNs and a very small set of CNs within a control facility at any time due to the full transfer of control as a plane moves across sectors of airspace. The set of CNs may be assumed to all share the same network prefix. These CNs are also involved in other data flows over the same access network that the MR is attached to, managing other flights within the sector. These CNs are often geographically and topologically much closer to the MR in comparison to a single fixed HA.

2.1.1.2. Airline Operational Services Domain

Data flows for Airline Operational Services (AOS) are not critical to the safety of the passengers or aircraft, but are needed for the business operations of airlines operating flights, and may affect the

profitability of an airline's flights. Most of these data flows are sourced by MNNs that are part of the flight management system or sensor nodes on an aircraft, and are terminated at CNs located near an airline's headquarters or operations center. AOS traffic may include detailed electronic passenger manifests, passenger ticketing and rebooking traffic, and complete electronic baggage manifests. When suitable bandwidth is available (currently on the surface when connected to a wired Gatelink system), "airplane health information" data transfers of between 10 and several-hundred megabytes of data are likely, and in the future, it is expected that the In-Flight Entertainment (IFE) systems may receive movie refreshes of data running into the gigabyte range.

Currently, these flows are often short messages that record the timing of events of a flight, engine performance data, etc., but may be longer flows that upload weather or other supplementary data to an aircraft. In addition, email-like interactive messaging may be used at any time during a flight. For instance, messages can be exchanged before landing to arrange for arrival-gate services to be available for handicapped passengers, refueling, food and beverage stocking, and other needs. This messaging is not limited to landing preparation, though, and may occur at any stage of flight.

The equipment comprising these MNNs and CNs has similar considerations to the equipment used for the ATS domain. A key difference between ATS and AOS is that AOS data flows are routed to CNs that may be much more geographically remote to the aircraft than CNs used by ATS flows, as AOS CNs will probably be located at an airline's corporate data center or headquarters. The AOS CNs will also probably be static for the lifetime of the flight, rather than dynamic like the ATS CNs. An HA used for AOS may be fairly close topologically to the CNs, and RO may not be as big of a benefit for AOS, since simple event logging is more typical than time-critical interactive messaging. For the small number of messaging flows, however, the CNs are geographically (but not necessarily topologically) very close to the aircraft. Additionally, since AOS communication is more advisory in nature than ATS, rather than safety-critical, AOS flows are less sensitive to tunnel inefficiencies than ATS flows. For these reasons, in this document, we consider AOS data flow concerns with RO mechanisms to not be full requirements, but instead consider them under the desirable properties in [Section 4](#).

[2.1.3](#). Passenger Services Domain

The MNNs involved in the Passenger Information and Entertainment Services (PIES) domain are mostly beyond the direct control of any single authority. The majority of these MNNs are VMNs and personal

property brought onboard by passengers for the duration of a flight, and thus it is unreasonable to assume that they be preloaded with special software or operating systems. These MNNs run stock Internet applications like web browsing, email, and file transfer, often through VPN tunnels. The MNNs themselves are portable electronics such as laptop computers and mobile smartphones capable of connecting to an onboard wireless access network (e.g. using 802.11). To these MNN devices and users, connecting to the onboard network is identical to connecting to any other terrestrial "hotspot" or typical wireless LAN. The MNNs are completely oblivious to the fact that this access network is on an airplane and possibly moving around the globe. The users are not always technically-proficient and may not be capable of performing any special configuration of their MNNs or applications.

A small number of PIES MNNs may be LFNs that store and distribute cached media content (e.g. movies and music), or may provide gaming services to passengers. Due to the great size of the data stored on these LFNs compared to the anemic bandwidth available air-to-ground, these LFNs will probably not attempt to communicate off-board at all during the course of a flight, but will wait to update their content via either high-speed links are available on the ground, or via removable media inserted by the flight crew. Any data flow needed for billing passengers for access to content can probably also be performed after landing. Since it seems like these LFNs are unlikely to require off-board communications during the course of a flight, we do not consider them in this document.

The PIES domain is not critical to safety-of-life, but is merely an added comfort or business service to passengers. Since PIES applications may consume much more bandwidth than the available links used in other domains, the PIES MNNs may have their packets routed through a separate high-bandwidth link, that is not used by the ATS data flows. For instance, several service providers are planning to offer passenger Internet access during flight at DSL-like rates, just as the Connexion by Boeing system did. Several airlines also plan to offer onboard cellular service to their passengers, possibly utilizing Voice-over-IP for transport. Due to the lack of criticality and the likelihood of being treated independently, in this document, PIES MNN concerns are not considered as input to requirements in [Section 3](#). The R0 solution should be optimized for ATS and AOS needs, and consider PIES as a secondary concern.

With this in consideration, the PIES domain is also the most likely to utilize NEMO for communications in the near-term since relatively little regulations and beaurocracy are involved in deploying new technology in this domain, and IP-based PIES systems have previously been developed and deployed (although not using NEMO) [9]. For these reasons, PIES concerns factor heavily into the

desirable properties in [Section 4](#) outside of the mandatory requirements.

Some PIES nodes are currently using 2.5G/3G links for mobile data services, and these may be able to migrate to an IP-based onboard mobile network, when available.

[2.2.](#) Space Exploration Scenarios

This section describes some features of the network environments found in space exploration that are relevant to selecting an appropriate NEMO RO mechanism. It should be noted that IPv4-based mobile routing has been demonstrated onboard the UK-DMC satellite and that the documentation on this serves as a useful reference for understanding some of the goals and configuration issues for certain types of space use of NEMO [\[10\]](#). This section assumes space use of NEMO within the "near-Earth" range of space (i.e. not for communications between the Earth and Mars, or other "deep-space" locations). Note, that NEMO is currently being considered for use out to Lunar distances. No strong distinction is made here between civilian versus military use, or exploration mission versus Earth-observing or other mission types; our focus is on civilian exploration missions, but we believe that many of the same basic concerns are relevant to these other mission types.

In space communications, a high degree of bandwidth asymmetry is often present, with the uplink from the ground to a craft typically being multiple orders of magnitude slower than the downlink from the craft to the ground. This means that the RO overhead may be negligible on the downlink but significant for the uplink. An RO scheme that minimizes the amount of signaling from CNs to an MN is desirable, since these uplinks may be low-bandwidth to begin with (possibly only several kilobits per second). Since the uplink is used for sending commands, it should not be blocked for long periods while serializing long RO signaling packets; any RO signaling from the CN to MNs must not involve large packets.

For unmanned space flight, the MNs onboard a spacecraft consist almost entirely of LFN sensing devices and processing devices that send telemetry and science data to CNs on the ground and actuator devices that are commanded from the ground in order to control the craft. Robotic lunar rovers may serve as VMNs behind an MR located on a lander or orbiter, but these rovers will contain many independent instruments and could probably be configured as an MR and LFNs instead of using a single VMN address.

It can be assumed that for manned spaceflight, at least, multiple MRs will be present and on-line simultaneously for fast failover. These

will usually be multihomed over space links in diverse frequency bands, and so multiple-prefixes can be expected, especially since some links will be direct to ground stations while others may be bent-pipe repeated through satellite relays like TDRSS. For unmanned missions, if low weight and power are more critical, it is likely that only a single MR and single link/prefix may be present.

In some modes of spacecraft operation, all communications may go through a single onboard computer (or a Command and Data Handling system as on the International Space Station) rather than directly to the MNNs themselves, so there is only ever one MNN behind an MR that is in direct contact with off-board CNs. In this case removing the MR and using simple host-based Mobile IPv6 rather than NEMO is possible, but an MR is more desirable because it could be part of a modular communications adapter that is used in multiple diverse missions to bridge on-board buses and intelligently manage space links, rather than re-creating these capabilities per-mission if using simple Mobile IPv6 with a single Command and Data Handling node that varies widely between spacecraft. Also, all visions for the future involve network-centric operations where the direct addressability and accessibility of end devices and data is crucial. As network-centric operations become more prevalent, application of NEMO is likely to be needed to increase the flexibility of data flow.

The MRs and MNNs onboard a spacecraft are highly-customized computing platforms, and adding custom code or complex configurations in order to obtain NEMO RO capabilities is feasible, although it should not be assumed that any amount of code or configuration maintenance is possible after launch. The RO scheme as it is initially configured should continue to function throughout the lifetime of an asset.

For manned space flight, additional MNNs on spacesuits and astronauts may be present and used for applications like two-way voice conversation or video-downlink. These MNNs could be reusable and reconfigured per-flight for different craft or mission network designs, but it is still desirable for them to be able to autoconfigure themselves and they may move between nested or non-nested MRs during a mission. For instance, if astronauts move between two docked spacecraft, each craft may have its own local MR and wireless coverage that the suit MNNs will have to reconfigure for. It is desirable if a RO solution can respond appropriately to this change in locality, and not cause high levels of packet loss during the transitional period. It is also likely that these MNNs will be part of Personal Area Networks (PANs), and so may appear either directly as MNNs behind the main MR onboard, or might have their own MR within the PAN and thus create a nested (or even multi-level nested) NEMO configuration.

3. Required Characteristics

This section lists requirements that specify the absolute minimal technical and/or functional properties that a NEMO RO mechanism must possess to be usable for aeronautical and space communications.

In the recent work done by the International Civil Aviation Organization (ICAO) to identify viable mobility technologies for providing IP services to aircraft, a set of technical criteria was developed [4] [11]. The nine required characteristics listed in this document can be seen as directly descended from these ICAO criteria, except here we have made them much more specific and focused for the NEMO technology. The original ICAO criteria were more general and used for comparing the features of different mobility solutions (e.g. mobility techniques based on routing protocols versus transport protocols versus Mobile IP, etc.). Within the text describing each requirement in this section, we provide the high-level ICAO criteria from which it evolved.

These requirements for aeronautics are generally similar to or in excess of the requirements for space exploration, so we do not add any additional requirements specifically for space exploration. In addition, the lack of a standards body regulating performance and safety requirements for space exploration means that the requirements for aviation are much easier to agree upon and base within existing requirements frameworks. After consideration, we believe that the set of aviation-based requirements outlined here also fully suffices for space exploration.

3.1. Req1 - Separability

An RO scheme MUST have the ability to be bypassed by traffic types that desire to use bidirectional tunnels through an HA.

3.1.1. Rationale for Aeronautics - Separability

For the aeronautics community there are basically three types of traffic (classes of service): "critical traffic" used for air traffic control and safety of flight, "signaling traffic" used for air operations management, and "user traffic" used by passenger services.

For critical and signaling traffic, there may be a concern with taking the path selected by a RO mechanism, since this may be considered a less trustworthy or less reliable path for some reason than the MRHA tunnel that represents a more-known quantity, or for capacity properties or regulatory reasons. This is not to say that the RO-selected path really is less safe to use, just that it may be perceived as such due to a lack of prior probing or unknown

information about its delay or capacity properties. Since the MRHA tunnel's path has known properties, it may be preferred over an unknown RO-selected path for certain restricted classes of data flows (e.g. ATS or spacecraft command and control).

For this reason, an RO scheme must have the ability to be bypassed by applications that desire to use bidirectional tunnels through an HA. This desire could be expressed through a policy database similar to the Security Policy Database used by IPsec, for instance, but the specific means of signaling or configuring the expression of this desire by applications is left as a detail for the specific RO specifications.

This requirement was derived from ICAO's TC-1 - "The approach should provide a means to define data communications that can be carried only over authorized paths for the traffic type and category specified by the user."

3.2. Req2 - Multihoming

Req2 - RO schemes MUST permit an MR to be simultaneously connected to multiple access networks, having multiple prefixes and Care-Of Addresses in a MONAMI6 context.

3.2.1. Rationale for Aeronautics - Multihoming

Multiple factors drive a requirement for multihoming capabilities. For ATS safety-of-life critical traffic, the need for high availability suggests a basic multihoming requirement. The regulatory and operational difficulty in deploying new systems and transitioning away from old ones also implies that a mix of access technologies may be in use at any given time, and may require simultaneous use. Another factor is that the multiple domains of applications onboard may actually be restricted in what data links they are allowed to use based on regulations and policy, so at certain times or locations, PIES data flows may have to use distinct access links from those used by ATS data flows.

This drives the requirement that an RO solution MUST allow for an MR to be connected to multiple access networks simultaneously and have multiple CoAs in use simultaneously. The selection of a proper CoA and access link to use per-packet may be either within or outside the scope of the RO solution. As a minimum, if an RO solution is integrable with the MONAMI6 basic extensions, and does not preclude their use, then this requirement can be considered to be satisfied.

It is not this requirement's intention that an RO scheme itself provide multihoming, but rather simply to exclude RO techniques whose

use is not possible in multihomed scenarios.

This requirement was derived from ICAO's TC-2 - "The approach should enable an aircraft to both roam between and to be simultaneously connected to multiple independent air-ground networks."

3.3. Req3 - Latency

Req3 - An RO solution MUST be capable of configuring itself (and reconfiguring after mobility events) without blocking unoptimized packet flow during its initiation and before or after transitions in the active access links.

3.3.1. Rationale for Aeronautics - Latency

It is possible that an RO scheme may take longer to setup or involve more signaling than the basic NEMO MRHA tunnel maintenance that occurs during an update to the MR's active CoAs when the set of usable access links changes. During this period of flux, it may be important for applications to be able to immediately get packets onto the ground network, especially considering that connectivity may have been blocked for some period of time while link-layer and NEMO procedures for dealing with the transition occurred. Also, when an application starts for the first time, the RO scheme may not have previous knowledge related to the CN and may need to perform some setup before an optimized path is available. If the RO scheme blocks packets either through queueing or dropping while it is configuring itself, this could result in unacceptable delays.

Thus, when transitions in the MR's set of active access links occurs, the RO scheme MUST NOT block packets from using the MRHA tunnel if the RO scheme requires more time to setup or configure itself than the basic NEMO tunnel maintenance. Additionally, when an application flow is started, the RO scheme MUST allow packets to immediately be sent, perhaps without the full benefit of RO, if the RO scheme requires additional time to configure a more optimal path to the CN.

This requirement was derived from ICAO's TC-3 - "The approach should minimize latency during establishment of initial paths to an aircraft, during handoff, and during transfer of individual data packets."

3.4. Req4 - Availability

Req4 - An RO solution MUST NOT imply a single point of failure, whether that be a single MR, a single HA, or other point within the ground network.

3.4.1. Rationale for Aeronautics - Availability

A need for high availability of connectivity to ground networks arises from the use of IP networking for carrying safety-of-life critical traffic. For this reason, single points of failure need to be avoided. If an R0 solution assumes either a single MR onboard, a single HA, or some similar vulnerable point, and is not usable when redundant elements are present and in use then it will not be acceptable. An R0 solution also MUST NOT itself imply a single point of failure.

This does not mention specific redundancy mechanisms for MRs, HAs, or other networking elements, so as long as some reasonable method for making each component redundant fits within the assumptions of the R0 mechanism, this requirement can be considered satisfied.

There is no intention to support "Internet-less" operation through this requirement. When an MR is completely disconnected from the majority of the network it is intended to communicate, including its HA, there is no requirement for it to be able to retain any communications involving parties outside the mobile networks managed by itself.

This requirement was derived from ICAO's TC-4 - "The approach should have high availability which includes not having a single point of failure."

3.5. Req5 - Packet Loss

Req5 - An R0 scheme MUST NOT cause packets to be dropped at any point in operation, when they would not normally have been dropped in a non-R0 configuration. (Integrity)

3.5.1. Rationale for Aeronautics - Packet Loss

It is possible that some R0 schemes could cause data packets to be lost during transitions in R0 state or due to unforeseen packet filters along the R0-selected path. This could be difficult for an application to detect and respond to in time. For this reason, an R0 scheme MUST NOT cause packets to be dropped at any point in operation, when they would not normally have been dropped in a non-R0 configuration. If duplicating a small number of packets sent over both the MRHA tunnel and the optimized path is possible until the optimized path can be verified to function for a particular data flow, then this could be sufficient to meet this requirement (it is safe to assume that the applications are robust to this duplication).

This requirement does not necessarily imply make-before-break in

transitioning between links. The intention is that during the handoff period, the R0 scheme itself should not produce losses that would not have occurred if R0 had been disabled.

This requirement was derived from ICAO's TC-5 - "The approach should not negatively impact end-to-end data integrity, for example, by introducing packet loss during path establishment, handoff, or data transfer."

It is understood that this may be a requirement that is not easily implementable with regards to R0. Furthermore Req1, Separability, may be sufficient. In addition, the ability for a R0 implementation to determine the appropriate path for particular systems is probably out of scope. Regardless, this requirement is currently being retained in order to ensure it gets addressed in some manner via either R0 or elsewhere. This requirement is likely to be removed, but we do not wish to lose the thought at this time.

3.6. Req6 - Scalability

Req6 - An R0 scheme MUST be simultaneously usable by ten thousand nodes without overloading the ground network or routing system.

3.6.1. Rationale for Aeronautics - Scalability

Several thousand aircraft may be in operation at some time, each with perhaps several hundred MNNs onboard. The number of active spacecraft using IP will be multiple orders of magnitude smaller than this over at least the next decade, so the aeronautical needs are more stringent in terms of scalability to large numbers of MRs. It would be a non-starter if the combined use of an R0 technique by all of the MRs in the network caused ground networks provisioned within the realm of typical long-haul private telecommunications networks (like the FAA's Telecommunications Infrastructure (FTI) or NASA's NISN) to be overloaded or melt-down under the R0 signaling load or amount of rapid path changes for multiple data flows.

Thus, an R0 scheme MUST be simultaneously usable by ten thousand nodes without overloading the ground network or routing system.

This requirement was derived from ICAO's TC-6 - "The approach should be scaleable to accommodate anticipated levels of aircraft equipage."

3.7. Req7 - Efficient Signaling

An R0 scheme MUST be capable of efficient signaling

3.7.1. Rationale for Aeronautics - Efficient Signaling

The amount of bandwidth available for aeronautical and space communications has historically been quite small in comparison to the desired bandwidth (e.g. in the case of VDL links the bandwidth is 8 kbps of shared resources). This situation is expected to persist for at least several more years. Links tend to be provisioned based on estimates of application needs (which could well prove wrong if either demand or the applications in-use themselves do not follow expectations), and do not leave much room for additional networking protocol overhead. Since every byte of available air-ground link capacity that is used by signaling for NEMO RO is likely to delay bytes of application data and reduce application throughput, it is important that the NEMO RO scheme's signaling overhead scales up much more slowly than the throughput of the flows RO is being performed on. This way as higher-rate datalinks are deployed along with more bandwidth-hungry applications, the NEMO RO scheme will be able to safely be discounted in capacity planning.

This requirement was derived from ICAO's TC-7 - "The approach should result in throughput which accommodates anticipated levels of aircraft equipage."

3.8. Req8 - Security

Req8 - IPsec MUST be usable over the RO scheme, and the data used to make RO decisions MUST be authenticable, perhaps using some form of IPsec.

3.8.1. Rationale for Aeronautics - Security

Security based on IPsec is widely understood and modules qualified for use in aeronautical and space exploration environments are currently available off-the-shelf due to the US Department of Defense HAIPE work. Other security frameworks may not be as attractive due to a lack of familiarity or available flight-qualified systems. If an RO solution precludes the use of IPsec by applications, this would make it undesirable or unusable. Likewise, since the IPsec standards seem to be the only widely-agreed upon framework for implementing packet authentication and other security services, the use of IPsec to authenticate or otherwise secure RO signaling and other data used in the RO process is also desirable, but other reasonable authentication mechanisms may be acceptable for RO signalling.

These issues motivate the requirement that IPsec MUST be usable by applications utilizing the RO scheme, and the data used to make RO decisions MAY be authenticable using IPsec.

This requirement was extrapolated from ICAO's TC-8 - "The approach should be secure."

3.9. Req9 - Adaptability

Req9 - New applications, potentially using new transport protocols or IP options MUST be possible within an R0 scheme.

3.9.1. Rationale for Aeronautics - Adaptability

The concepts of operations are not fully developed for network-centric command and control and other uses of IP-based networks in aeronautical and space environments. The exact application protocols, data flow characteristics, and even transport protocols that will be used in either transitional and final operational concepts are not completely defined yet, and may even change with deployment experience. If the R0 solution assumes that some amount of preconfiguration of flow properties (port number ranges, etc.) is required, this could make the integration of new applications or protocols difficult. An R0 scheme might assume this in order to classify between flows that prefer the MRHA tunnel to an optimized path per requirement Req1, for instance, among other reasons. Since flag days or other such large-scale synchronized configuration updates are to be avoided to ensure high availability, the R0 scheme MUST NOT fail on the use of unexpected higher layer protocols (transport and above).

This drives the requirement that new applications, potentially using new transport protocols must be usable within an R0 scheme, and MUST NOT involve global reconfiguration events for MRs.

This requirement was derived from ICAO's TC-9 - "The approach should be scalable to accommodate anticipated transition to new IP-based communication protocols."

4. Desirable Characteristics

In this section, we identify some of the properties of the system that are not strict requirements due to either being difficult to quantify or to being features that are not immediately needed, but may provide additional benefits that would help encourage adoption.

4.1. Des1 - Configuration

For ATS systems, complex configurations are known to increase uncertainty in context, human error and the potential for undesirable (unsafe) states to be reached [[12](#)]. Since R0 alters the

communications context between an MNN and CN, 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.

For CNs at large airports, the Binding Cache state management functions may be simultaneously dealing with hundreds of airplanes with multiple service providers, and a volume of mobility events due to arrivals and departures. The ability to have simple interfaces for humans to access the Binding Cache configuration and alter it in case of errors are desirable, if this does not interfere with the RO protocol mechanisms themselves.

4.2. Des2 - Nesting

It is desirable if the RO mechanism supports RO for nested MRs, since it is possible that for PIES and astronaut spacesuits, that PANs with MRs will need to be supported. For oceanic flight, ATS and AOS may also benefit from the capability of nesting MRs between multiple planes to provide a "reachback" to terrestrial groundstations rather than relying solely on lower rate HF or satellite systems. In either case, this mode of operation is beyond current strict requirements and is merely desirable.

4.3. Des3 - System Impact

Low complexity in systems engineering and configuration management is desirable in building and maintaining systems using the RO mechanism. This property may be difficult to quantify, judge, and compare between different RO techniques, but a mechanism that is perceived to have lower impact on the complexity of the network communications system should be favored over an otherwise equivalent mechanism (with regards to the requirements listed above). This is somewhat different than Des1 (Configuration), in that Des1 refers to operation and maintenance of the system once deployed, whereas Des3 is concerned with the initial design, deployment, transition, and later upgrade path of the system.

4.4. Des4 - VMN Support

At least LFNs and LMNs MUST be supported by a viable RO solution for aeronautics, as these local nodes are within the ATS and AOS domains. VMNs within the PIES domain are expected to be numerous, and it is strongly desirable for them to be supported by the RO technique, but not strictly required.

4.5. Des5 - Generality

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. For instance, an RO solution that is usable within VANETs [\[13\]](#) or consumer electronics equipment [\[14\]](#) could satisfy this. The goal is for the technology to be more widely-used and maintained outside the relatively-small aeronautical networking community and its vendors, in order to make acquisitions and training faster, easier, and cheaper. This could also allow aeronautical networking to possibly benefit from future RO scheme optimizations and developments whose research and development is funded and performed externally by the broader industry and academic communities.

5. Security Considerations

This document does not create any security concerns in and of itself. The security properties of any NEMO RO scheme that is to be used in aeronautics and space exploration are probably much more stringent than for more general NEMO use, due to the safety-of-life and/or national security issues involved. The required security properties are described under Req8 of [Section 3](#) within this document.

When deploying in operational networks where network layer security may be mandated (e.g. virtual private networks), the interaction between this and NEMO RO techniques should be carefully considered to ensure that the security mechanisms do not undo the route optimization by forcing packets through a less-optimal overlay or underlay. For instance, when IPsec tunnel use is required, the locations of the tunnel endpoints can force sub-optimal end-to-end paths to be taken.

6. IANA Considerations

This document neither creates nor updates any IANA registries.

7. Acknowledgments

Input from several parties is indirectly included in this document. Participants in the MPI mailing list and BoF efforts helped to shape the document. The NEMO and MONAMI6 working group participants were instrumental in completing this document. Specific suggestions from Steve Bretmersky, Thierry Ernst, Tony Li, Jari Arkko, Phillip Watson, Roberto Baldessari, Carlos Jesus Bernardos Cano, and Eivan Cerasi

were incorporated into this document.

Wesley Eddy's work on this document was performed at NASA's Glenn Research Center, while in support of NASA's Advanced Communications Navigations and Surveillance Architectures and System Technologies (ACAST) project, and the NASA Space Communications Architecture Working Group (SCAWG).

8. Changes from [draft-eddy-nemo-aero-reqs-01](#)

Note, this section will be removed upon publication as an RFC.

At the IETF69 nemo meeting, consensus was reached to have the aeronautics, car-to-car, electronics devices, and any other groups develop individual requirements documents. At a later date, these individual requirements documents are likely to be combined into one document. Thus, the [draft-eddy-nemo-aero-req-01](#) document has been revised as follows:

The document has been renamed as [draft-eddy-nemo-aero-reqs-02](#). Once MEXT gets organized, this document will become an ietf-mext working document per IETF69 consensus. However, in order to get this document out for comment during the transition period, it has been released at this time as an individual submission.

In Req1 - Separability the wording has been changed such that selection is based on traffic types rather than applications. It was noted during the IETF69 NEMO meeting that requiring a R0 to determine separability by application is not practical but use of techniques such as traffic classification may be.

Req5 - Integrity has been changed to Req5 - Packet Loss. Within the IETF integrity has a different meaning than packet loss. Within the aviation community, integrity, as used here, refers to path integrity or data flow integrity. Within the IETF, integrity implies good data rather than good path. Thus, in order to improve clarity within the IETF, the wording has been changed.

Req7 - Throughput has been changed to Req7 - Efficient Signaling. Overall throughput is relative to the bandwidth of the available links and is relatively independent of the R0 implementation. Rather, the real requirement is for R0 to be deployable over low bandwidth systems (e.g. VHDL links of 8 kbps for today's aeronautical communication links). Thus, the wording has been changed to reflect the actual requirement.

Acronyms have been defined the first time they are used.

The requirements section has been re-written to read more like a contract. Thus, we have requirements and rationale with the rationale being numbered paragraphs. The idea is that, if so desired at a later date, this should enable combining of the various requirement documents into a single requirements document. In addition, we hope that the changes will make this document easier for an implementer to utilize.

9. Informative References

- [1] Ernst, T. and H. Lach, "Network Mobility Support Terminology", [draft-ietf-nemo-terminology-06](#) (work in progress), November 2006.
- [2] Ernst, T., "Network Mobility Support Goals and Requirements", [draft-ietf-nemo-requirements-06](#) (work in progress), November 2006.
- [3] Ivancic, W., "Multi-Domained, Multi-Homed Mobile Networks", [draft-ivancic-mobile-platforms-problem-00](#) (work in progress), September 2006.
- [4] Davis, T., "Mobile Internet Platform Aviation Requirements", [draft-davis-mip-aviationreq-00](#) (work in progress), September 2006.
- [5] Devarapalli, V., Wakikawa, R., Petrescu, A., and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol", [RFC 3963](#), January 2005.
- [6] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", [RFC 3775](#), June 2004.
- [7] Ng, C., "Network Mobility Route Optimization Problem Statement", [draft-ietf-nemo-ro-problem-statement-03](#) (work in progress), September 2006.
- [8] ICAO Asia/Pacific Regional Office, "Required Communication Performance (RCP) Concepts - An Introduction", Informal South Pacific ATS Coordinating Group 20th meeting, Agenda Item 7, January 2006.
- [9] Dul, A., "Global IP Network Mobility", presentation at IETF 62 plenary , March 2005.

- [10] Ivancic, W., Paulsen, P., Stewart, D., Shell, D., Wood, L., Jackson, C., Hodgson, D., Northam, J., Bean, N., Miller, E., Graves, M., and L. Kurisaki, "Secure, Network-centric Operations of a Space-based Asset: Cisco Router in Low Earth Orbit (CLEO) and Virtual Mission Operations Center (VMOC)", NASA Technical Memorandum TM-2005-213556, May 2005.
- [11] ICAO WG-N SWG1, "Analysis of Candidate ATN IPS Mobility Solutions", Meeting #12, Working Paper 6, Bangkok, Thailand, January 2007.
- [12] ICAO, "Threat and Error Management (TEM) in Air Traffic Control", ICAO Preliminary Edition, October 2005.
- [13] Baldessari, R., "C2C-C Consortium Requirements for NEMO Route Optimization", [draft-baldessari-c2ccc-nemo-req-01](#) (work in progress), July 2007.
- [14] Ng, C., "Consumer Electronics Requirements for Network Mobility Route Optimization", [draft-ng-nemo-ce-req-00](#) (work in progress), July 2007.
- [15] CCSDS, "Cislunar Space Internetworking: Architecture", CCSDS 000.0-G-1 Draft Green Book, December 2006.
- [16] NASA Space Communication Architecture Working Group, "NASA Space Communication and Navigation Architecture Recommendations for 2005-2030", SCAWG Final Report, May 2006.
- [17] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

[Appendix A](#). Basics of IP-based Aeronautical Networking

The current standards for aeronautical networking are based on the ISO OSI networking stack and are referred to as the Aeronautical Telecommunications Network (ATN). While standardized, the ATN has not been fully deployed and seems to be in only limited use compared to its full vision and potential. The International Civil Aviation Organization (ICAO) is a part of the United Nations that produces standards for aeronautical communications. The ICAO has recognized that an ATN based on OSI lacks the widespread commercial network support required for the successful deployment of new more bandwidth-intensive ATN applications, and has recently been working towards a new IP-based version of the ATN.

Supporting mobility in an IP-based network may be vastly different

than it is in the OSI-based ATN which uses the IDRP interdomain routing protocol to recompute routing tables as mobile networks change topological points of attachment. ICAO recognizes this and has studied various mobility techniques based on link, network, transport, routing, and application protocols [[11](#)].

Work done within ICAO has identified the NEMO technology as a promising candidate for use in supporting global IP-based mobile networking. The main concerns with NEMO have been with its current lack of route optimization support and its potentially complex configuration requirements in a large airport environment with multiple service providers and 25 or more airlines sharing the same infrastructure.

[Appendix B](#). Basics of IP-based Space Networking

IP itself is only in limited operational use for communicating with spacecraft currently (e.g. the SSTL DMC satellites are one example). Future communications architectures include IP-based networking as an essential building-block, however. The Consultative Committee for Space Data Systems (CCSDS) has a working group that is producing a network architecture for using IP-based communications in both manned and unmanned near-Earth missions, and has international participation towards this goal [[15](#)]. NASA's Space Communications Architecture Working Group (SCAWG) also has developed an IP-based multi-mission networking architecture [[16](#)]. Neither of these is explicitly based on Mobile IP technologies, but NEMO is usable within these architectures, and they may be extended to include NEMO when/if the need becomes apparent.

Authors' Addresses

Wesley M. Eddy
Verizon Federal Network Systems
NASA Glenn Research Center
21000 Brookpark Road, MS 54-5
Cleveland, OH 44135
USA

Email: weddy@grc.nasa.gov

Will Ivancic
NASA Glenn Research Center
21000 Brookpark Road, MS 54-5
Cleveland, OH 44135
USA

Phone: +1-216-433-3494
Email: William.D.Ivancic@grc.nasa.gov

Terry Davis
Boeing Commercial Airplanes
P.O.Box 3707 MC 07-25
Seattle, WA 98124-2207
USA

Phone: 206-280-3715
Email: Terry.L.Davis@boeing.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

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.

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).

