        Network Mobility Route Optimization Requirements for
  Operational Use in Aeronautics and Space Exploration Mobile Networks

Abstract

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

   Substantial input to these requirements was given by aeronautical
   communications experts outside the IETF, including members of the
   International Civil Aviation Organization (ICAO) and other
   aeronautical communications standards bodies.

Status of This Memo

   This memo provides information for the Internet community. It does
   not specify an Internet standard of any kind. Distribution of this
   memo is unlimited.

Copyright Notice

Table of Contents

1. Introduction

   As background, the Network Mobility (NEMO) terminology and NEMO goals
   and requirements documents are suggested reading ([4], [5]).

   The base NEMO standard [1] extends Mobile IPv6 [2] for singular
   mobile hosts in order to be used by Mobile Routers (MRs) supporting
   entire mobile networks. 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 MR and a
   NEMO-supporting Home Agent (HA) placed at some relatively stable
   point in the network. NEMO does not provide Mobile IPv6's Route
   Optimization (RO) features to Mobile Network Nodes (MNNs) other than
   to the NEMO MR itself. Corresponding Nodes (CNs) that communicate

with MNNs behind an MR do so through the HA and the bidirectional
Mobile Router - Home Agent (MRHA) tunnel. Since the use of this
tunnel may have significant drawbacks [6], 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 on board 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 must often use distinct air-ground links and
service providers.

For aeronautics, the main disadvantage of using NEMO bidirectional
tunnels is that airlines operate flights that traverse multiple
continents, and a single plane may fly around the entire world over a
span of a couple days. If a plane uses a static HA on a single
continent, then for a large percentage of the time, when the plane is
not on the same continent as the HA, a great amount of delay is
imposed by using the MRHA tunnel. Avoiding the delay from
unnecessarily forcing packets across multiple continents is the
primary goal of pursuing NEMO RO for aeronautics.

Other properties of the aeronautics and space environments amplify
the known issues with NEMO bidirectional MRHA tunnels [6] 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 Automatic Repeat Request (ARQ) mechanisms and
        underprovisioned radio links. Furthermore, for space
        exploration and for aeronautical communications systems that
        pass through geosynchronous satellites, 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, mean 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 [7].

  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:
     RFC 4888 [6] identifies fragmentation due to encapsulation as
     an artifact of tunneling. While links used in the aeronautics
     and space domains are error-prone and may cause loss of
     fragments on the initial/final hop(s), considerations for
     fragmentation along the entire tunneled path are the same as
     for the terrestrial domain.

  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 [3].
Although this document does not specify an actual protocol, but
rather specifies 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, this section describes some
   operational characteristics to explain how access networks, HAs, and
   CNs are configured and distributed geographically and topologically
   in aeronautical and space network architectures. This may be useful
   in determining which classes of RO techniques within the known
   solution space [8] are feasible.

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. local networks available while on a gate) 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. These handovers may result in
   having a different access router and a change in CoA, though the use
   of local mobility management (e.g., [9]) may limit the changes in CoA
   to only handovers between different providers or types of data links.

   The characteristics of a data flow between a CN and MNN varies both
   depending on the data flow's domain and on the particular application
   within the domain. Even within the three aeronautical domains
   described below, there are varying classes of service that are
   regulated differently (e.g., for emergencies versus nominal
   operations), 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. Air Traffic Services Domain

   The MNNs involved in Air Traffic Services (ATS) consist of pieces of
   avionics hardware on board an aircraft that are 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 multihomed MR that is at least multi-interfaced and possibly
   multi-prefixed as well, in NEMO terminology.

   The existing ATS link technologies may be too anemic for a complete
   IP-based ATS communications architecture (link technologies and
   acronyms are briefly defined in Appendix A). At the time of this
   writing, the ICAO is pursuing future data link standards that support
   higher data rates. Part of the problem is limited spectrum, pursued
   under ICAO ACP-WG-F, "Spectrum Management", and part of the problem
   is the data link protocols themselves, pursued under ICAO ACP-WG-T,
   "Future Communications Technology". ACP-WG-T has received inputs
   from studies on a number of potential data link protocols, including
   B-AMC, AMACS, P34, LDL, WCDMA, and others. Different link
   technologies may be used in different stages of flight, for instance
   802.16 in the surface and terminal area, P34 or LDL en route, and
   satcom in oceanic flight. Both current and planned data links used
   for Passenger Information and Entertainment Services (PIES) and/or
   Airline Operational Services (AOS), such as the satcom links employed
   by passenger Internet-access systems, support much higher data rates
   than current ATS links.

   Since, for ATS, the MRs and MNNs are under regulatory control and are
   actively tested and maintained, it is not completely unreasonable to
   assume that special patches or software be run on these devices to
   enable NEMO RO. In fact, since these devices are accessed by skilled
   technicians and professionals, it may be that some special
   configuration is required for NEMO RO. Of course, simplicity in set
   up and configuration is highly preferable, however, and the desirable
   feature labeled "Des1" later in this document prefers solutions with
   lower configuration state and overhead. To minimize costs of

ownership and operations, it is also highly desirable for only widely
available, off-the-shelf operating systems or network stacks to be
required, but this is not a full requirement.

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 and take place 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 be topologically close to one another. 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.

The MNNs and CNs used for ATS will support IP services, as IP is the
basis of the new Aeronautical Telecommunications Network (ATN)
architecture being defined by ICAO. Some current ATS ground systems
run typical operating systems, like Solaris, Linux, and Windows, on
typical workstation computer hardware. There is some possibility for
an RO solution to require minor changes to these CNs, though it is
much more desirable if completely off-the-shelf CN machines and
operating systems can be used. Later in this document, the security
requirements suggest that RO might be performed with mobility anchors
that are topologically close to the CNs, rather than directly to CNs
themselves. This could possibly mean that CN modifications are not
required.

During the course of a flight, there are several events for which an
RO solution should consider the performance implications:

o Initial session creation with an ATS CN (called "Data Link Logon"
   in the aeronautical jargon).

o Transfer of control between ATS CNs, resulting in regional
   differences in where the controlling CN is located.

o Aircraft-initiated contact with a non-controlling ATS CN, which
   may be located anywhere, without relation to the controlling CN.

o Non-controlling, ATS, CN-initiated contact with the aircraft.

o   Aircraft transition between one access link to another, resulting
    in change of CoA.

o   Concurrent use of multiple access links with different care-of
    addresses.

2.1.2. 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 link at a gate), "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
(e.g., television programming or recent news updates) running into
the multi-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, though this depends on how

applications are written -- whether they use centralized servers or
exchange messages directly. 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 desirable properties, which are discussed in Section 4.

Future AOS MNNs and CNs can be expected to implement IPv6 and conform
to the new IPv6-based ATN Standards and Recommended Practices (SARPS)
that ICAO is defining. AOS CNs have similar hardware and software
properties as described for ATS above.

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 on board 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.

The largest class of PIES CNs consists of typical web servers and
other nodes on the public Internet. It is not reasonable to assume
that these can be modified specifically to support a NEMO RO scheme.
Presently, these CNs would be mostly IPv4-based, though an increasing
number of IPv6 PIES CNs are expected in the future. This document
does not consider the problem of IPv4-IPv6 transition, beyond the
assumption that either MNNs and CNs are running IPv6 or a transition
mechanism exists somewhere within the network.

A small number of PIES MNNs may be LFNs that store and distribute
cached media content (e.g., movies and music) or that 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 available on the ground or

removable media inserted by the flight crew. However, if a higher
bandwidth link were affordably available, it might be used in-flight
for these purposes, but supporting this is not a requirement. Data
flows needed for billing passengers for access to content are
relatively low bandwidth and are currently done in-flight. The
requirements of these data flows are less stringent than those of
ATS, however, so they are not specifically considered here.

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 former 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 RO 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 bureaucracy are involved in deploying new
technology in this domain and since IP-based PIES systems have
previously been developed and deployed (although not using NEMO)
[10]. 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 on board 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 [11]. 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 MNNs must not involve large packets.

For unmanned space flight, the MNNs on board 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 online simultaneously for fast failover. These
will usually be multihomed over space links in diverse frequency
bands, and so multiple access network prefixes can be expected to be
in use simultaneously, especially since some links will be direct to
ground stations while others may be bent-pipe repeated through
satellite relays like the Tracking and Data Relay Satellite System
(TDRSS). This conforms to the (n,1,1) or (n,n,1) NEMO multihoming
scenarios [12]. 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, conforming to the (1,1,1) or (1,n,1) NEMO
multihoming scenarios [12].

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. However, an MR is more desirable because it could be part
of a modular communications adapter that is used in multiple diverse

missions to bridge onboard buses and intelligently manage space
links. This is cheaper and leads to faster development time 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 on board 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 spacecrafts, each craft may have its own local MR
and wireless coverage that the suit MNNs will have to reconfigure
for. It is desirable if an 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 on board or 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 ([13], [14]). 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 and the problem of RO within NEMO.

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.

It is understood that different solutions may be needed for
supporting different domains. This may mean either different NEMO RO
solutions or different mobility solutions entirely. Divergent
solutions amongst the domains are acceptable, though preferably
avoided if possible.

An underlying requirement that would be assumed by the use of Mobile
IP technology for managing mobility (rather than a higher-layer
approach) is that IP addresses used both within the mobile network
and by CNs to start new sessions with nodes within the mobile network
remain constant throughout the course of flights and operations. For
ATS and AOS, this allows the Home Addresses (HoAs) to serve as node
identifiers, rather than just locators, and for PIES it allows common
persistent applications (e.g., Voice over IP (VoIP) clients, VPN
clients, etc.) to remain connected throughout a flight. Prior
aeronautical network systems like the prior OSI-based ATN and
Connexion by Boeing set a precedent for keeping a fixed Mobile
Network Prefix (MNP), though they relied on interdomain routing
protocols (IDRP and BGP) to accomplish this, rather than NEMO
technology. This requirement applies to the selection in general of
a mobility management technology, and not specifically to an RO
solution once NEMO has been decided on for mobility management.

3.1. Req1 - Separability

Since RO may be inappropriate for some flows, an RO scheme MUST
support configuration by a per-domain, dynamic RO policy database.
Entries in this database can be similar to those used in IPsec
security policy databases in order to specify either bypassing or
utilizing RO for specific flows.

3.1.1. Rationale for Aeronautics - Separability

   Even if RO is available to increase the performance of a mobile
   network's traffic, it may not be appropriate for all flows.

   There may also be a desire to push certain flows through the MRHA
   path, rather than performing RO, to enable them to be easily recorded
   by a central service.

   For these reasons, 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.

   In addition, it is expected that the use of NEMO technology be
   decided on a per-domain basis, so that it is possible that, for some
   domains, separate MRs or even non-NEMO mobility techniques are used.
   This requirement for an RO policy database only applies to domains
   that utilize NEMO.

   This requirement was derived from ICAO's TC-1 [15] - "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."

   One suggested approach to traffic separation is multi-addressing of
   the onboard networks, with treatment of a traffic domain determined
   by the packet addresses used. However, there are other techniques
   possible for meeting this requirement, and so multi-addressing is not
   itself a requirement. The Req1 requirement we describe above is
   intended for separating the traffic within a domain that makes use of
   NEMO based on flow properties (e.g., short messaging flows vs. longer
   file transfers or voice flows).

3.2. Req2 - Multihoming

   An RO solution MUST support an MR having multiple interfaces and MUST
   allow a given domain to be bound to a specific interface. It MUST be
   possible to use different MNPs for different domains.

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 on board may actually be restricted in what data links
they are allowed to use, based on regulations and policy; thus, 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 (i.e., registration of
multiple CoAs and flow bindings) 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.

In terms of NEMO multihoming scenarios [12], it MUST be possible to
support at least the (n,1,n) and (n,n,n) 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

While an RO solution is in the process of setting up or
reconfiguring, packets of specified flows MUST be capable of using
the MRHA tunnel.

3.3.1. Rationale for Aeronautics - Latency

It is possible that an RO scheme may take longer to set up 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 set

up 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 set up 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

An RO solution MUST be compatible with network redundancy mechanisms
and MUST NOT prevent fallback to the MRHA tunnel if an element in an
optimized path fails.

An RO mechanism MUST NOT add any new single point of failure for
communications in general.

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 RO solution assumes either a single onboard MR, a
single HA, or some similar vulnerable point, and is not usable when
the network includes standard reliability mechanisms for routers,
then the RO technique will not be acceptable. An RO solution also
MUST NOT itself imply a single point of failure.

This requirement specifies that the RO solution itself does not
create any great new fragility. Although in basic Mobile IPv6 and
NEMO deployments, the use of a single HA implies a single point of
failure, there are mechanisms enabling the redundancy of HAs (e.g.,
[16]). It is assumed that some HA-redundancy techniques would be
employed to increase robustness in an aeronautical setting. It
should also be understood that the use of RO techniques decreases
dependence on HAs in the infrastructure and allows a certain level of
robustness to HA failures in that established sessions using RO may

be able to operate based on Binding Cache entries even after an HA
failure. With RO, an HA failure primarily impacts the ability to
connect new application flows to a mobile network.

If a failure occurs in a path selected by an RO technique, then that
RO technique MUST NOT prevent fallback to the MRHA path for affected
traffic.

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 RO
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 with which 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

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

3.5.1. Rationale for Aeronautics - Packet Loss

It is possible that some RO schemes could cause data packets to be
lost during transitions in RO state or due to unforeseen packet
filters along the RO-selected path. This could be difficult for an
application to detect and respond to in time. For this reason, an RO
scheme SHOULD NOT cause packets to be dropped at any point in
operation, when they would not normally have been dropped in a non-RO
configuration.

As an attempt at optimizing against packet loss, some techniques may,
for some time, duplicate packets sent over both the MRHA tunnel and
the optimized path. If this results in duplicate packets being
delivered to the application, this is also unacceptable.

This requirement does not necessarily imply make-before-break in transitioning between links. The intention is that during the handoff period, the RO scheme itself should not produce losses (or duplicates) that would not have occurred if RO 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 RO. Furthermore Req1, Separability, may be sufficient in allowing loss-sensitive and duplicate-sensitive flows to take the MRHA path.

3.6. Req6 - Scalability

An RO scheme MUST be simultaneously usable by the MNNs on hundreds of thousands of craft without overloading the ground network or routing system. This explicitly forbids injection of BGP routes into the global Internet for purposes of RO.

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 RO 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 the NASA Integrated Services Network (NISN)) to be overloaded or melt-down under the RO signaling load or amount of rapid path changes for multiple data flows.

Thus, an RO scheme MUST be simultaneously usable by the MNNs on hundreds of thousands of craft without overloading the ground network or routing system. The scheme must also be tolerant to the delay and/or loss of initial packets, which may become more pervasive in future Internet routing and addressing architectures [17].

Since at least one traffic domain (PIES) requires connectivity to the Internet and it is possible that the Internet would provide transport for other domains at some distant point in the future, this requirement explicitly forbids the use of techniques that are known to scale poorly in terms of their global effects, like BGP, for the

purposes of RO. The previous OSI-based ATN system used IDRP and an
"island" concept for maintaining connectivity to the mobile network
but was not tested on a large scale deployment. The Connexion by
Boeing system used BGP announces and withdrawals as a plane moved
across the globe in order to maintain connectivity [10]. This was
found to contribute to a significant amount of churn in the global
Internet routing tables, which is undesirable for a number of
reasons, and must be avoided in the future.

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

The specific scaling factor for the number of aircraft used in our
version of the requirement is an order of magnitude larger than the
estimated equipage cited in an ICAO draft letter-of-intent to ARIN
for an IPv6 prefix allocation request. There were several other
estimates that different groups had made, and it was felt in the IETF
that using a larger estimate was more conservative. It should be
noted that even with this difference of an order of magnitude, the
raw number is still several orders of magnitude lower than that of
estimated cellular telephone users, which might use the same protocol
enhancements as the cellular industry has also adopted Mobile IP
standards.

3.7. Req7 - Efficient Signaling

An RO scheme MUST be capable of efficient signaling in terms of both
size and number of individual signaling messages and the ensemble of
signaling messages that may simultaneously be triggered by concurrent
flows.

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 data links are deployed along with more
bandwidth-hungry applications, the NEMO RO scheme will be able to
safely be discounted in capacity planning.

Note that in meeting this requirement, an RO technique must be
efficient in both the size and number of individual messages that it
sends, as well in the ensemble of messages sent at one time (for
instance, to give RO to multiple ongoing flows following a handover),
in order to prevent storms of packets related to RO.

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

   For the ATS/AOS domains, there are three security sub-requirements:

   1. The RO scheme MUST NOT further expose MNPs on the wireless link
       than already is the case for NEMO basic support.

   2. The RO scheme MUST permit the receiver of a binding update (BU)
       to validate an MR's ownership of the CoAs claimed by an MR.

   3. 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.

3.8.1. Rationale for Aeronautics - Security

   The security needs are fairly similar between ATS and AOS, but vary
   widely between the ATS/AOS domains and PIES. For PIES, the traffic
   flows are typical of terrestrial Internet use and the security
   requirements for RO are identical to those of conventional Mobile
   IPv6 RO. For ATS/AOS, however, there are somewhat more strict
   requirements, along with some safe assumptions that designers of RO
   schemes can make. Below, we describe each of these ATS/AOS issues,
   but do not further discuss PIES RO security.

   The first security requirement is driven by concerns expressed by ATS
   communications engineers. The concern is driven by current air-
   ground links to a craft and their lack of security, which has allowed
   eavesdroppers to track individual flights in detail. Protecting the
   MNP from exposure has been expressed as a requirement by this
   community, though the security of the RO system should not depend on

secrecy of the MNP. The RO scheme should use some reasonable
security mechanisms in order to both protect RO signaling via strong
authentication and encrypt the MNP from being visible over air-ground
links.

The second security requirement is driven by the risk of flooding
attacks that are started by an attacker redirecting an MNP's traffic
to some target victim CoA. To protect bindings to bogus CoAs from
being sent, the RO scheme must somehow validate that an MR actually
possesses any CoAs that it claims. For the purposes of aeronautics,
it is safe to assume ingress filtering is in place in the access
networks.

To protect against "rogue" MRs or abuse of compromised MRs, the RO
scheme MUST be capable of checking that an MR is actually authorized
to perform a binding update for a specific MNP. To meet this
requirement, it can be assumed that some aeronautical organization
authority exists who can provide the required authorization, possibly
in the form of a certificate that the MR possesses, signed by the
aeronautical authority.

It is also reasonable to assume trust relationships between each MR
and a number of mobility anchor points topologically near to its CNs
(these anchor points may be owned by the service providers), but it
is not reasonable to assume that trust relationships can be
established between an MR and any given CN itself. Within the
onboard networks for ATS and AOS, it is reasonable to assume that the
LFNs and MRs have some trust relationship.

It is felt by many individuals that by the time the IP-based ATN
grows into production use, there will be a global ATN-specific Public
Key Infrastructure (PKI) usable for ATS, though it is agreed that
such a PKI does not currently exist and will take time to develop
both technically and politically. This PKI could permit the
establishment of trust relationships among any pair of ATS MNNs, MRs,
or CNs through certificate paths, in contrast to the more limited
amount of trust relationships described in the previous paragraph.
While it has been suggested that early test and demonstration
deployments with a more limited-scale PKI deployment can be used in
the near-term, as a global PKI is developed, some parties still feel
that assuming a global PKI may be overly bold in comparison to
assuming trust relationships with anchor points. It is always
possible to scale the anchor point assumption up if a PKI develops
that allows the CNs themselves to become the anchor points. It is
not possible to go back down in the other direction if a global PKI
never emerges.

This requirement was extrapolated from ICAO's TC-8 - "The approach
should be secure" and made more specific with help from the MEXT
working group.

3.9. Req9 - Adaptability

Applications using new transport protocols, IPsec, or new IP options
MUST be possible within an RO 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 or final operational
concepts are not completely defined yet, and may even change with
deployment experience. The RO solution itself should allow all
higher-layer protocols, ports, and options to be used.

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
that 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 reaching
undesirable (unsafe) states [18]. Since RO 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 is 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, 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 ground stations 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. It is also noted that there are other ways
   to support these communications scenarios using routing protocols or
   other means outside of NEMO.

   Loop-detection, in support of nesting, is specifically not a
   requirement at this stage of ATN and space network designs, due to
   both the expectation that the operational environments are carefully
   controlled and inherently avoid loops and the understanding that
   scenarios involving nesting are not envisioned in the near future.

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 MUST be supported by a viable RO solution for
   aeronautics, as these local nodes are within the ATS and AOS domains.
   If Mobile IPv6 becomes a popular technology used by portable consumer
   devices, 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. LMNs are potentially present
   in future space exploration scenarios, such as manned exploration
   missions to the moon and Mars.

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 Vehicular ad hoc Networks (VANETs) [19] or consumer
   electronics equipment [20] 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.

   Under an assumption of closed and secure backbone networks, the air-
   ground link is the weakest portion of the network and most
   susceptible to injection of packets, flooding, and other attacks.
   Future air-ground data links that will use IP are being developed
   with link-layer security as a concern. This development can assist
   in meeting one of this document's listed security requirements (that
   MNPs not be exposed on the wireless link), but the other requirements
   affect the RO technology more directly without regard to the presence
   or absence of air-ground link-layer security.

   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. Acknowledgments

   Input from several parties is indirectly included in this document.
   Participants in the Mobile Platform Internet (MPI) mailing list and
   BoF efforts helped to shape the document, and the early content was

borrowed from MPI problem statement and proposed requirements
documents ([21], [13]). The NEMO and MONAMI6 working group
participants were instrumental in completing this document. The
participants in the MEXT interim meeting February 7th and 8th of 2008
in Madrid were critical in solidifying these requirements. Specific
suggestions from Steve Bretmersky, Thierry Ernst, Tony Li, Jari
Arkko, Phillip Watson, Roberto Baldessari, Carlos Jesus Bernardos
Cano, Eivan Cerasi, Marcelo Bagnulo, Serkan Ayaz, Christian Bauer,
Fred Templin, Alexandru Petrescu, Tom Henderson, and Tony Whyman were
incorporated into this document.

Wesley Eddy's work on this document was performed at NASA's Glenn
Research Center, primarily 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) in 2005 and 2006.

7. References

7.1. Normative References

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

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

[3]   Bradner, S., "Key words for use in RFCs to Indicate Requirement
        Levels", BCP 14, RFC 2119, March 1997.

7.2. Informative References

[4]   Ernst, T. and H-Y. Lach, "Network Mobility Support
        Terminology", RFC 4885, July 2007.

[5]   Ernst, T., "Network Mobility Support Goals and Requirements",
        RFC 4886, July 2007.

[6]   Ng, C., Thubert, P., Watari, M., and F. Zhao, "Network Mobility
        Route Optimization Problem Statement", RFC 4888, July 2007.

[7]   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.

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

   [9]   Kempf, J., "Goals for Network-Based Localized Mobility
         Management (NETLMM)", RFC 4831, April 2007.

   [10] Dul, A., "Global IP Network Mobility", Presentation at IETF
         62 Plenary, March 2005.

   [11] 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.

   [12] Ng, C., Ernst, T., Paik, E., and M. Bagnulo, "Analysis of
         Multihoming in Network Mobility Support", RFC 4980,
         October 2007.

   [13] Davis, T., "Mobile Internet Platform Aviation Requirements",
         Work in Progress, September 2006.

   [14] ICAO WG-N SWG1, "Analysis of Candidate ATN IPS Mobility
         Solutions", Meeting #12, Working Paper 6, Bangkok, Thailand,
         January 2007.

   [15] Davis, T., "Aviation Global Internet Operations Requirements",
         ICAO WG-N, Sub-Working-Group N1, Information Paper #4 (IP4),
         September 2006.

   [16] Wakikawa, R., "Home Agent Reliability Protocol", Work
         in Progress, July 2009.

   [17] Zhang, L. and S. Brim, "A Taxonomy for New Routing and
         Addressing Architecture Designs", Work in Progress, March 2008.

   [18] ICAO, "Threat and Error Management (TEM) in Air Traffic
         Control", ICAO Preliminary Edition, October 2005.

   [19] Baldessari, R., "C2C-C Consortium Requirements for NEMO Route
         Optimization", Work in Progress, July 2007.

   [20] Ng, C., "Consumer Electronics Requirements for Network Mobility
         Route Optimization", Work in Progress, February 2008.

   [21] Ivancic, W., "Multi-Domained, Multi-Homed Mobile Networks",
         Work in Progress, September 2006.

   [22] CCSDS, "Cislunar Space Internetworking: Architecture", CCCSDS
         000.0-G-1 Draft Green Book, December 2006.

   [23] NASA Space Communication Architecture Working Group, "NASA
         Space Communication and Navigation Architecture Recommendations
         for 2005-2030", SCAWG Final Report, May 2006.

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 IPv6-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 Inter-Domain Routing
   Protocol (IDRP) 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 [14].

   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.

   A significant challenge to the deployment of networking technologies
   to aeronautical users is the low capability of existing air-ground
   data links for carrying IP-based (or other) network traffic. Due to
   barriers of spectrum and certification, production of new standards
   and equipment for the lower layers below IP is slow. Currently
   operating technologies may have data rates measured in the several
   kbps range, and it is clear that supporting advanced IP-based
   applications will require new link technologies to be developed
   simultaneously with the development of networking technologies
   appropriate for aeronautics.

   In addition to well-known commercial data links that can be adapted
   for aeronautical use, such as Wideband Code-Division Multiple Access
   (WCDMA) standards or the IEEE 802.16 standard, several more
   specialized technologies either exist or have been proposed for air-
   ground use:

o VHF Data Link (VDL) specifies four modes of operation in the
   117.975 – 137 MHz range that are capable of supporting different
   mixes of digital voice and data at fairly low rates. The low
   rates are driven by the need to operate within 25 kHz channels
   internationally allocated for aeronautical use. VDL mode 2 is
   somewhat widely deployed on aircraft and two global service
   providers support VDL access networks. Experiences with VDL mode
   2 indicate that several kbps of capacity delivered to a craft can
   be expected in practice, and the use of long timers and a
   collision avoidance algorithm over a large physical space
   (designed to operate at 200 nautical miles) limit the performance
   of IP-based transport protocols and applications.

o Aircraft Communications and Reporting System (ACARS) is a
   messaging system that can be used over several types of underlying
   RF data links (e.g., VHF, HF, and satellite relay). ACARS
   messaging automates the sending and processing of several types of
   event notifications over the course of a flight. ACARS in general
   is a higher-level messaging system, whereas the more specific
   "Plain Old ACARS" (POA) refers to a particular legacy RF interface
   that the ACARS system employed prior to the adoption of VDL and
   other data links. Support for IP-based networking and advanced
   applications over POA is not feasible.

o Broadband Aeronautical Multi-carrier Communications (B-AMC) is a
   hybrid cellular system that uses multi-carrier CDMA from ground-
   to-air and Orthogonal Frequency Division Multiplexing (OFDM) in
   the air-to-ground direction. B-AMC runs in the L-band of spectrum
   and is adapted from the Broadband-VHF (B-VHF) technology
   originally developed to operate in the VHF spectrum. L-band use
   is intended to occupy the space formerly allocated for Distance
   Measuring Equipment (DME) using channels with greater bandwidth
   than are available than in the VHF band, where analog voice use
   will continue to be supported. B-AMC may permit substantially
   higher data rates than existing deployed air-ground links.

o All-Purpose Multi-Channel Aviation Communications System (AMACS)
   is an adaptation of the Global System for Mobile Communications
   (GSM) physical layer to operate in the L-band with 50 - 400 kHz
   channels and use VDL mode 4's media access technique. AMACS may
   permit data rates in the several hundred kbps range, depending on
   specific channelization policies deployed.

o Project 34 (P34) is a wideband public-safety radio system capable
   of being used in the L-band. P34 is designed to offer several
   hundred kbps of capacity specifically for IP-based packet
   networking. It uses OFDM in 50, 100, or 150 kHz channels and

exact performance will depend on the particular operating band,
range (guard time), and channelization plan configured in
deployment.

o L-Band Data Link (LDL) is another proposal using the L-band based
   on existing technologies. LDL adapts the VDL mode 3 access
   technique and is expected to be capable of up to 100 kbps.

Appendix B. Basics of IP-based Space Networking

IP itself is only in limited operational use for communicating with
spacecraft currently (e.g., the Surry Satellite Technology Limited
(SSTL) Disaster Monitoring Constellation (DMC) satellites). 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 [22]. NASA's Space Communications Architecture
Working Group (SCAWG) also has developed an IP-based multi-mission
networking architecture [23]. 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 0Y-96
   Seattle, WA 98124-2207
   USA

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