Network Working Group   F. Templin, Ed.
Internet-Draft   G. Saccone
Intended status: Informational   Boeing Research & Technology
Expires: February 17, 2019   G. Dawra
                                                             LinkedIn
                                                            A. Lindem
                                                            V. Moreno
                                                  Cisco Systems, Inc.
                                                      August 16, 2018

           A Simple BGP-based Mobile Routing System for the Aeronautical
                        Telecommunications Network
                        draft-templin-atn-bgp-08.txt

Abstract

   The International Civil Aviation Organization (ICAO) is investigating
   mobile routing solutions for a worldwide Aeronautical
   Telecommunications Network with Internet Protocol Services (ATN/IPS).
   The ATN/IPS will eventually replace existing communication services
   with an IPv6-based service supporting pervasive Air Traffic
   Management (ATM) for Air Traffic Controllers (ATC), Airline
   Operations Controllers (AOC), and all commercial aircraft worldwide.
   This informational document describes a simple and extensible mobile
   routing service based on industry-standard BGP to address the ATN/IPS
   requirements.

Copyright Notice

Table of Contents

1. Introduction

   The worldwide Air Traffic Management (ATM) system today uses a
   service known as Aeronautical Telecommunications Network based on
   Open Systems Interconnection (ATN/OSI). The service is used to
   augment controller to pilot voice communications with rudimentary
   short text command and control messages. The service has seen
   successful deployment in a limited set of worldwide ATM domains.

   The International Civil Aviation Organization [ICAO] is now
   undertaking the development of a next-generation replacement for ATN/
   OSI known as Aeronautical Telecommunications Network with Internet
   Protocol Services (ATN/IPS). ATN/IPS will eventually provide an
   IPv6-based [RFC8200] service supporting pervasive ATM for Air Traffic
   Controllers (ATC), Airline Operations Controllers (AOC), and all

commercial aircraft worldwide. As part of the ATN/IPS undertaking, a
new mobile routing service will be needed. This document presents an
approach based on the Border Gateway Protocol (BGP) [RFC4271].

Aircraft communicate via wireless aviation data links that typically
support much lower data rates than terrestrial wireless and wired-
line communications. For example, some Very High Frequency (VHF)-
based data links only support data rates on the order of 32Kbps and
an emerging L-Band data link that is expected to play a key role in
future aeronautical communications only supports rates on the order
of 1Mbps. Although satellite data links can provide much higher data
rates during optimal conditions, like any other aviation data link
they are subject to errors, delay, disruption, signal intermittence,
degradation due to atmospheric conditions, etc. The well-connected
ground domain ATN/IPS network should therefore treat each safety-of-
flight critical packet produced by (or destined to) an aircraft as a
precious commodity and strive for an optimized service that provides
the highest possible degree of reliability.

The ATN/IPS is an IPv6-based overlay network that assumes a worldwide
connected Internetworking underlay for carrying tunneled ATM
communications. The Internetworking underlay could be manifested as
a private collection of long-haul backbone links (e.g., fiber optics,
copper, SATCOM, etc.) interconnected by high-performance networking
gear such as bridges, switches, and routers. Such a private network
would need to connect all ATN/IPS participants worldwide, and could
therefore present a considerable cost for a large-scale deployment of
new infrastructure. Alternatively, the ATN/IPS could be deployed as
a secured overlay over the existing global public Internet. For
example, ATN/IPS nodes could be deployed as part of an SD-WAN or an
MPLS-WAN that rides over the public Internet via secured tunnels.
Further details of the Internetworking underlay design are out of
scope for this document.

The ATN/IPS further assumes that each aircraft will receive an IPv6
Mobile Network Prefix (MNP) that accompanies the aircraft wherever it
travels. ICAO is further proposing to assign each aircraft an entire
/56 MNP for numbering its on-board networks. ATCs and AOCs will
likewise receive IPv6 prefixes, but they would typically appear in
static (not mobile) deployments such as air traffic control towers,
airline headquarters, etc. Throughout the rest of this document, we
therefore use the term "MNP" when discussing an IPv6 prefix that is
delegated to any ATN/IPS end system, including ATCs, AOCs, and
aircraft. We also use the term Mobility Service Prefix (MSP) to
refer to an aggregated prefix assigned to the ATN/IPS by an Internet
assigned numbers authority, and from which all MNPs are delegated
(e.g., up to 2^{32} IPv6 /56 MNPs could be delegated from an IPv6 /24
MSP).

Connexion By Boeing [CBB] was an early aviation mobile routing
service based on dynamic updates in the global public Internet BGP
routing system. Practical experience with the approach has shown
that frequent injections and withdrawals of MNPs in the Internet
routing system can result in excessive BGP update messaging, slow
routing table convergence times, and extended outages when no route
is available. This is due to both conservative default BGP protocol
timing parameters (see Section 6) and the complex peering
interconnections of BGP routers within the global Internet
infrastructure. The situation is further exacerbated by frequent
aircraft mobility events that each result in BGP updates that must be
propagated to all BGP routers in the Internet that carry a full
routing table.

We therefore consider an approach using a BGP overlay network routing
system where a private BGP routing protocol instance is maintained
between ATN/IPS Autonomous System (AS) Border Routers (ASBRs). The
private BGP instance does not interact with the native BGP routing
system in the connected Internetworking underlay, and BGP updates are
unidirectional from "stub" ASBRs (s-ASBRs) to a small set of "core"
ASBRs (c-ASBRs) in a hub-and-spokes topology. No extensions to the
BGP protocol are necessary.

The s-ASBRs for each stub AS connect to a small number of c-ASBRs via
dedicated high speed links and/or tunnels across the Internetworking
underlay using industry-standard encapsulations (e.g., Generic
Routing Encapsulation (GRE) [RFC2784], IPsec [RFC4301], etc.). In
particular, tunneling must be used when neighboring ASBRs are
separated by many Internetworking underlay hops.

The s-ASBRs engage in external BGP (eBGP) peerings with their
respective c-ASBRs, and only maintain routing table entries for the
MNPs currently active within the stub AS. The s-ASBRs send BGP
updates for MNP injections or withdrawals to c-ASBRs but do not
receive any BGP updates from c-ASBRs. Instead, the s-ASBRs maintain
default routes with their c-ASBRs as the next hop, and therefore hold
only partial topology information.

The c-ASBRs connect to other c-ASBRs using internal BGP (iBGP)
peerings over which they collaboratively maintain a full routing
table for all active MNPs currently in service. Therefore, only the
c-ASBRs maintain a full BGP routing table and never send any BGP
updates to s-ASBRs. This simple routing model therefore greatly
reduces the number of BGP updates that need to be synchronized among
peers, and the number is reduced further still when intradomain
routing changes within stub ASes are processed within the AS instead
of being propagated to the core. BGP Route Reflectors (RRs)
[RFC4456] can also be used to support increased scaling properties.

The remainder of this document discusses the proposed BGP-based ATN/
IPS mobile routing service.

2. Terminology

The terms Autonomous System (AS) and Autonomous System Border Router
(ASBR) are the same as defined in [RFC4271].

The following terms are defined for the purposes of this document:

Air Traffic Management (ATM)
   The worldwide service for coordinating safe aviation operations.

Air Traffic Controller (ATC)
   A government agent responsible for coordinating with aircraft
   within a defined operational region via voice and/or data Command
   and Control messaging.

Airline Operations Controller (AOC)
   An airline agent responsible for tracking and coordinating with
   aircraft within their fleet.

Aeronautical Telecommunications Network with Internet Protocol
Services (ATN/IPS)
   A future aviation network for ATCs and AOCs to coordinate with all
   aircraft operating worldwide. The ATN/IPS will be an IPv6-based
   overlay network service that connects access networks via
   tunneling over an Internetworking underlay.

Internetworking underlay A connected wide-area network that supports
   overlay network tunneling and connects Radio Access Networks to
   the rest of the ATN/IPS.

Radio Access Network (RAN)
   An aviation radio data link service provider's network, including
   radio transmitters and receivers as well as supporting ground-
   domain infrastructure needed to convey a customer's data packets
   to the outside world. The term RAN is intended in the same spirit
   as for cellular operator networks and other radio-based Internet
   service provider networks. For simplicity, we also use the term
   RAN to refer to ground-domain networks that connect AOCs and ATCs
   without any aviation radio communications.

Core Autonomous System Border Router (c-ASBR) A BGP router located
   in the hub of the ATN/IPS hub-and-spokes overlay network topology.

Core Autonomous System The "hub" autonomous system maintained by all
   c-ASBRs.

   Stub Autonomous System Border Router (s-ASBR) A BGP router
      configured as a spoke in the ATN/IPS hub-and-spokes overlay
      network topology.

   Stub Autonomous System A logical grouping that includes all Clients
      currently associated with a given s-ASBR.

   Client An ATC, AOC or aircraft that connects to the ATN/IPS as a
      leaf node. The Client could be a singleton host, or a router that
      connects a mobile network.

   Proxy A node at the edge of a RAN that acts as an intermediary
      between Clients and s-ASBRs. From the Client's perspective, the
      Proxy presents the appearance that the Client is communicating
      directly with the s-ASBR. From the s-ASBR's perspective, the
      Proxy presents the appearance that the s-ASBR is communicating
      directly with the Client.

   Mobile Network Prefix (MNP) An IPv6 prefix that is delegated to any
      ATN/IPS end system, including ATCs, AOCs, and aircraft.

   Mobility Service Prefix (MSP) An aggregated prefix assigned to the
      ATN/IPS by an Internet assigned numbers authority, and from which
      all MNPs are delegated (e.g., up to 2**32 IPv6 /56 MNPs could be
      delegated from a /24 MSP).

3. ATN/IPS Routing System

   The proposed ATN/IPS routing system comprises a private BGP instance
   coordinated in an overlay network via tunnels between neighboring
   ASBRs over the Internetworking underlay. The overlay does not
   interact with the native BGP routing system in the connected
   underlying Internetwork, and each c-ASBR advertises only a small and
   unchanging set of MSPs into the Internetworking underlay routing
   system instead of the full dynamically changing set of MNPs. (For
   example, when the Internetworking underlay is the global public
   Internet the c-ASBRs advertise the MSPs in the public BGP Internet
   routing system.)

   In a reference deployment, one or more s-ASBRs connect each stub AS
   to the overlay using a shared stub AS Number (ASN). Each s-ASBR
   further uses eBGP to peer with one or more c-ASBRs. All c-ASBRs are
   members of the same core AS, and use a shared core ASN. Globally-
   unique public ASNs could be assigned, e.g., either according to the
   standard 16-bit ASN format or the 32-bit ASN scheme defined in
   [RFC6793].

The c-ASBRs use iBGP to maintain a synchronized consistent view of
all active MNPs currently in service. Figure 1 below represents the
reference deployment. (Note that the figure shows details for only
two s-ASBRs (s-ASBR1 and s-ASBR2) due to space constraints, but the
other s-ASBRs should be understood to have similar Stub AS, MNP and
eBGP peering arrangements.) The solution described in this document
is flexible enough to extend to these topologies.

```
    ...........................................................
    .                                                         .
    .               (:::)-.  <- Stub ASes ->  (:::)-.         .
    .   MNPs-> .-(:::::::::)          .-(:::::::::) <-MNPs     .
    .          `-(::::)-'              `-(::::)-'              .
    .           +-------+              +-------+              .
    .           |s-ASBR1+-----+    +-----+s-ASBR2|            .
    .           +--+----+ eBGP \   / eBGP +-----+-+           .
    .              \            \ /            /              .
    .               \eBGP        / \         /eBGP            .
    .                \          /   \        /                .
    .                +-------+   +-------+                     .
    .         eBGP+-----+c-ASBR |...|c-ASBR +-----+eBGP        .
    .  +-------+ /      +--+----+   +-----+-+      \ +-------+  .
    .  |s-ASBR +/       iBGP\   (:::)-.  /iBGP      \+s-ASBR |  .
    .  +-------+             .-(:::::::::)            +-------+  .
    .     .                .-(:::::::::::::::)-.         .      .
    .     .               (::::   Core AS   :::)        .      .
    .  +-------+           `-(:::::::::::::::)-'      +-------+  .
    .  |s-ASBR +\       iBGP/`-(:::::::-'\iBGP       /+s-ASBR |  .
    .  +-------+ \        +-+-----+   +----+--+     / +-------+  .
    .         eBGP+-----+c-ASBR |...|c-ASBR +-----+eBGP         .
    .               +-------+   +-------+                       .
    .              /                      \                    .
    .             /eBGP                    \eBGP               .
    .            /                          \                  .
    .         +---+---+              +-----+-+                 .
    .         |s-ASBR |              |s-ASBR |                 .
    .         +-------+              +-------+                 .
    .                                                         .
    .                                                         .
    .  <------------ Internetworking Underlay  -------------->  .
     ...........................................................
```
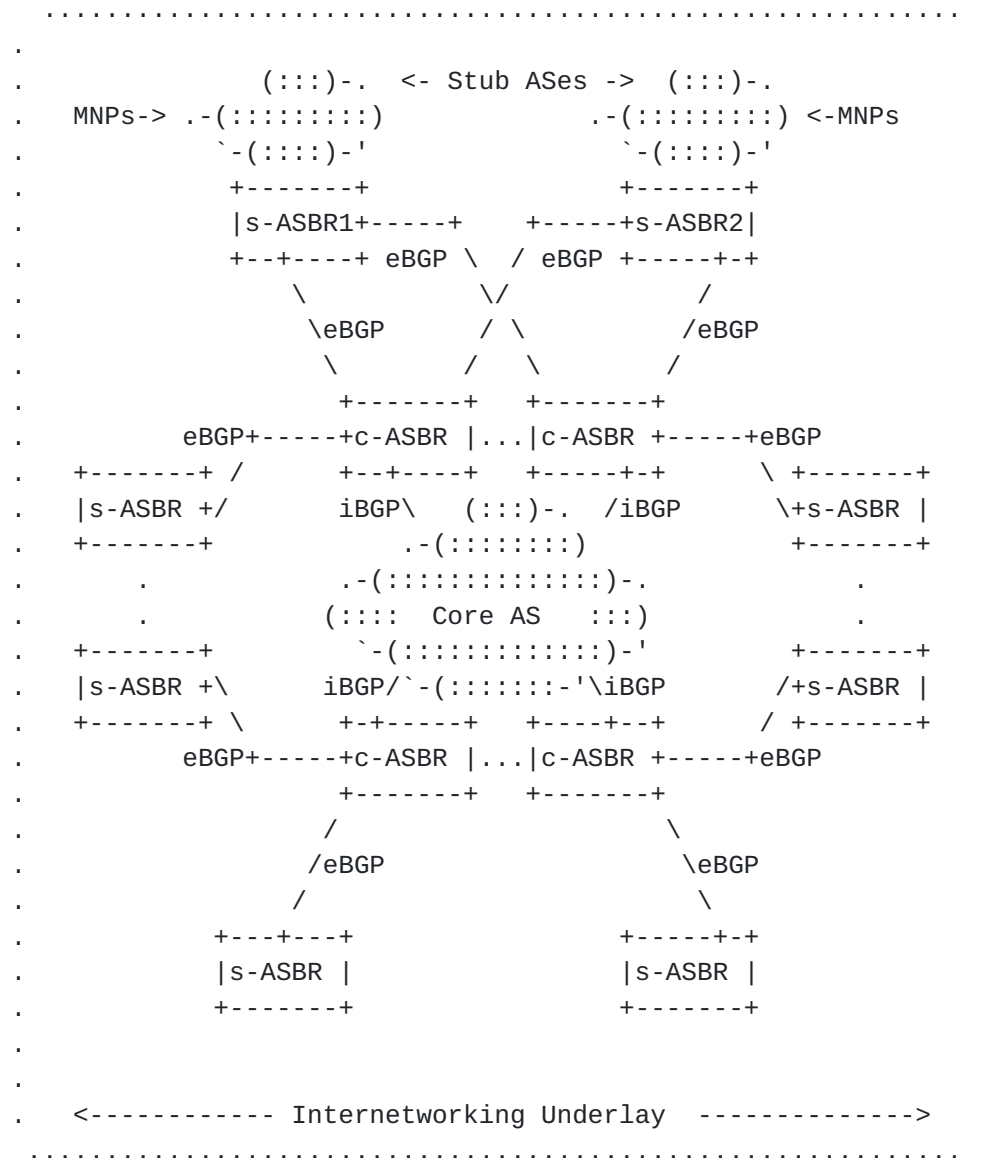
                    Figure 1: Reference Deployment

In the reference deployment, each s-ASBR maintains routes for active
MNPs that currently belong to its stub AS. In response to "Inter-
domain" mobility events, each s-ASBR will dynamically announces new
MNPs and withdraws departed MNPs in its eBGP updates to c-ASBRs.

Since ATN/IPS end systems are expected to remain within the same stub
AS for extended timeframes, however, intra-domain mobility events
(such as an aircraft handing off between cell towers) are handled
within the stub AS instead of being propagated as inter-domain eBGP
updates.

Each c-ASBR configures a black-hole route for each of its MSPs. By
black-holing the MSPs, the c-ASBR will maintain forwarding table
entries only for the MNPs that are currently active, and packets
destined to all other MNPs will correctly incur ICMPv6 Destination
Unreachable messages [RFC4443] due to the black hole route. (This is
the same behavior as for ordinary BGP routers in the Internet when
they receive packets for which there is no route available.) The
c-ASBRs do not send eBGP updates for MNPs to s-ASBRs, but instead
originate a default route. In this way, s-ASBRs have only partial
topology knowledge (i.e., they know only about the active MNPs
currently within their stub ASes) and they forward all other packets
to c-ASBRs which have full topology knowledge.

Scaling properties of this ATN/IPS routing system are limited by the
number of BGP routes that can be carried by the c-ASBRs. A 2015
study showed that BGP routers in the global public Internet at that
time carried more than 500K routes with linear growth and no signs of
router resource exhaustion [BGP]. A more recent network emulation
study also showed that a single c-ASBR can accommodate at least 1M
dynamically changing BGP routes even on a lightweight virtual
machine. Commercially-available high-performance dedicated router
hardware can support many millions of routes.

Therefore, assuming each c-ASBR can carry 1M or more routes, this
means that at least 1M ATN/IPS end system MNPs can be serviced by a
single set of c-ASBRs and that number could be further increased by
using RRs and/or more powerful routers. Another means of increasing
scale would be to assign a different set of c-ASBRs for each set of
MSPs. In that case, each s-ASBR still peers with one or more c-ASBRs
from each set of c-ASBRs, but the s-ASBR institutes route filters so
that it only sends BGP updates to the specific set of c-ASBRs that
aggregate the MSP. In this way, each set of c-ASBRs maintains
separate routing and forwarding tables so that scaling is distributed
across multiple c-ASBR sets instead of concentrated in a single
c-ASBR set. For example, a first c-ASBR set could aggregate an MSP
segment A::/32, a second set could aggregate B::/32, a third could
aggregate C::/32, etc. The union of all MSP segments would then
constitute the collective MSP(s) for the entire ATN/IPS.

In this way, each set of c-ASBRs services a specific set of MSPs that
they inject into the Internetworking underlay native routing system,
and each s-ASBR configures MSP-specific routes that list the correct

set of c-ASBRs as next hops. This design also allows for natural
incremental deployment, and can support initial medium-scale
deployments followed by dynamic deployment of additional ATN/IPS
infrastructure elements without disturbing the already-deployed base.
For example, a few more c-ASBRs could be added if the MNP service
demand ever outgrows the initial deployment.

4. ATN/IPS Radio Access Network (RAN) Model

Radio Access Networks (RANs) connect end system Clients such as
aircraft, ATCs, AOCs etc. to the ATN/IPS routing system. Clients may
connect to multiple RANs at once, for example, when they have both
satellite and cellular data links activated simultaneously. Clients
may further move between RANs in a manner that is perceived as a
network layer mobility event. Clients could therefore employ a
multilink/mobility routing service such as that discussed in
[I-D.templin-aerolink].

Clients register all of their active data link connections with their
serving s-ASBRs as discussed in Section 3. Clients may connect to
s-ASBRs either directly, or via a Proxy at the edge of the RAN.

Figure 2 shows the ATN/IPS RAN model where Clients connect to RANs
via aviation data links. Clients register their RAN addresses with a
nearby s-ASBR, where the registration process may be brokered by a
Proxy at the edge of the RAN.

```
            Data Link "A"      +--------+  Data Link "B"
               +----------- | Client |-----------+
              /               +--------+           \
             /                                      \
            /                                        \
        (:::)-.                                        (:::)-.
      .-(:::::::::) <- Radio Access Networks -> .-(:::::::::)
        `-(:::)-'                                  `-(:::)-'
        +-------+                                  +-------+
     ... | Proxy |  .......................... | Proxy |  ...
     .   +-------+                              +-------+   .
     .      ^^                                     ^^      .
     .      ||                                     ||      .
     .      ||               +--------+            ||      .
     .    ++============> | s-ASBR | <============++      .
     .                      +--------+                     .
     .                          | eBGP                     .
     .                        (:::)-.                       .
     .                      .-(:::::::::)                    .
     .                    .-(::: ATN/IPS :::)-.              .
     .                   (::::: BGP Routing ::::)            .
     .                    `-(:: System ::::)-'               .
     .                       `-(:::::::::-'                  .
     .                                                       .
     .                                                       .
     .   <------------- Internetworking Underlay ------------->   .
        ...........................................................
```

                  Figure 2: ATN/IPS RAN Architecture

   When a Client logs into a RAN, it specifies a nearby s-ASBR that it
   has selected to connect to the ATN/IPS. The login process is
   brokered by a Proxy at the border of the RAN, which then conveys the
   connection request to the s-ASBR via tunneling across the
   Internetworking underlay. The s-ASBR then registers the address of
   the Proxy as the address for the Client, and the Proxy forwards the
   s-ASBR's reply to the Client. If the Client connects to multiple
   RANs, the s-ASBR will register the addresses of all Proxies as
   addresses through which the Client can be reached.

   The s-ASBR represents all of its active Clients as MNP routes in the
   ATN/IPS BGP routing system. The s-ASBR's stub AS therefore consists
   of the set of all of its active Clients (i.e., the stub AS is a
   logical construct and not a physical construct). The s-ASBR injects
   the MNPs of its active Clients and withdraws the MNPs of its departed
   Clients via BGP updates to c-ASBRs. Since Clients are expected to
   remain associated with their current s-ASBR for extended periods, the
   level of MNP injections and withdrawals in the BGP routing system

will be on the order of the numbers of network joins, leaves and
s-ASBR handovers for aircraft operations (see: Section 6). It is
important to observe that fine-grained events such as Client mobility
and Quality of Service (QoS) signaling are coordinated only by
Proxies and the Client's current s-ASBRs, and do not involve other
ASBRs in the routing system. In this way, intradomain routing
changes within the stub AS are not propagated into the rest of the
ATN/IPS BGP routing system.

5. ATN/IPS Route Optimization

ATN/IPS end systems will frequently need to communicate with
correspondents associated with other s-ASBRs. In the BGP peering
topology discussed in Section 3, this can initially only be
accommodated by including multiple tunnel segments in the forwarding
path. In many cases, it would be desirable to eliminate extraneous
tunnel segments from this "dogleg" route so that packets can traverse
a minimum number of tunneling hops across the Internetworking
underlay. ATN/IPS end systems could therefore employ a route
optimization service such as that discussed in
[I-D.templin-aerolink].

A route optimization example is shown in Figure 3 and Figure 4 below.
In the first figure, multiple tunneled segments between Proxys and
ASBRs are necessary to convey packets between Clients associated with
different s-ASBRs. In the second figure, the optimized route tunnels
packets directly between Proxys without involving the ASBRs.

```
         +---------+                          +---------+
         | Client1 |                          | Client2 |
         +---v-----+                          +-----^---+
             *                                    *
             *                                    *
          (:::)-.                              (:::)-.
       .-(:::::::::::) <- Radio Access Networks -> .-(:::::::::::)
        `-(:::::)-'                             `-(:::::)-'
         +--------+                            +--------+
     ... | Proxy1 |  ........................  | Proxy2 |  ...
      .  +--------+                            +--------+   .
      .       **                                  **        .
      .        **                                  **       .
      .         **                                  **      .
      .          **                                  **     .
      .       +---------+                  +---------+       .
      .       | s-ASBR1 |                  | s-ASBR2 |       .
      .       +--+------+                  +-----+---+       .
      .           \   **      Dogleg      **   /            .
      .       eBGP\  **      Route      **   /eBGP          .
      .            \  **==============**   /                .
      .             +---------+   +---------+               .
      .             | c-ASBR1 |   | c-ASBR2 |               .
      .             +---+-----+   +----+----+               .
      .                 +--------------+                    .
      .                       iBGP                          .
      .                                                     .
      .   <------------- Internetworking Underlay ------------->   .
       .........................................................
```
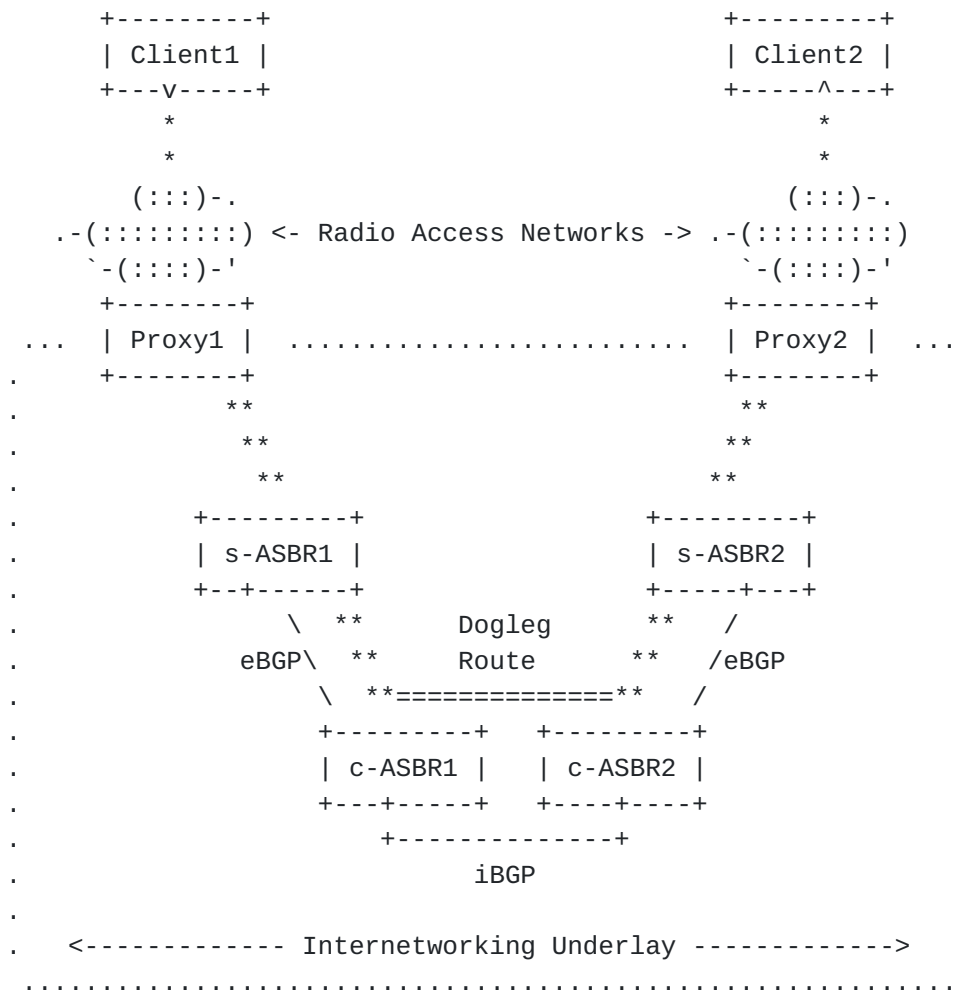
                  Figure 3: Dogleg Route Before Optimization

```
       +---------+                            +---------+
       | Client1 |                            | Client2 |
       +---v-----+                            +-----^---+
           *                                        *
           *                                        *
         (:::)-.                                  (:::)-.
      .-(:::::::::::)  <- Radio Access Networks -> .-(:::::::::)
        `-(:::::)-'                                  `-(:::::)-'
       +--------+                               +--------+
   ... | Proxy1 |  ........................      | Proxy2 |  ...
    .   +------v-+                               +--^-----+    .
    .          *                                      *        .
    .         *==============================*                 .
    .                                                          .
    .                                                          .
    .        +---------+                   +---------+         .
    .        | s-ASBR1 |                   | s-ASBR2 |         .
    .        +--+------+                   +-----+---+         .
    .            \                              /             .
    .        eBGP\                          /eBGP             .
    .             \                        /                  .
    .              +---------+   +---------+                   .
    .              | c-ASBR1 |   | c-ASBR2 |                   .
    .              +---+-----+   +----+----+                   .
    .                  +--------------+                        .
    .                        iBGP                             .
    .                                                          .
    .   <------------- Internetworking Underlay ------------->    .
     ..........................................................
```
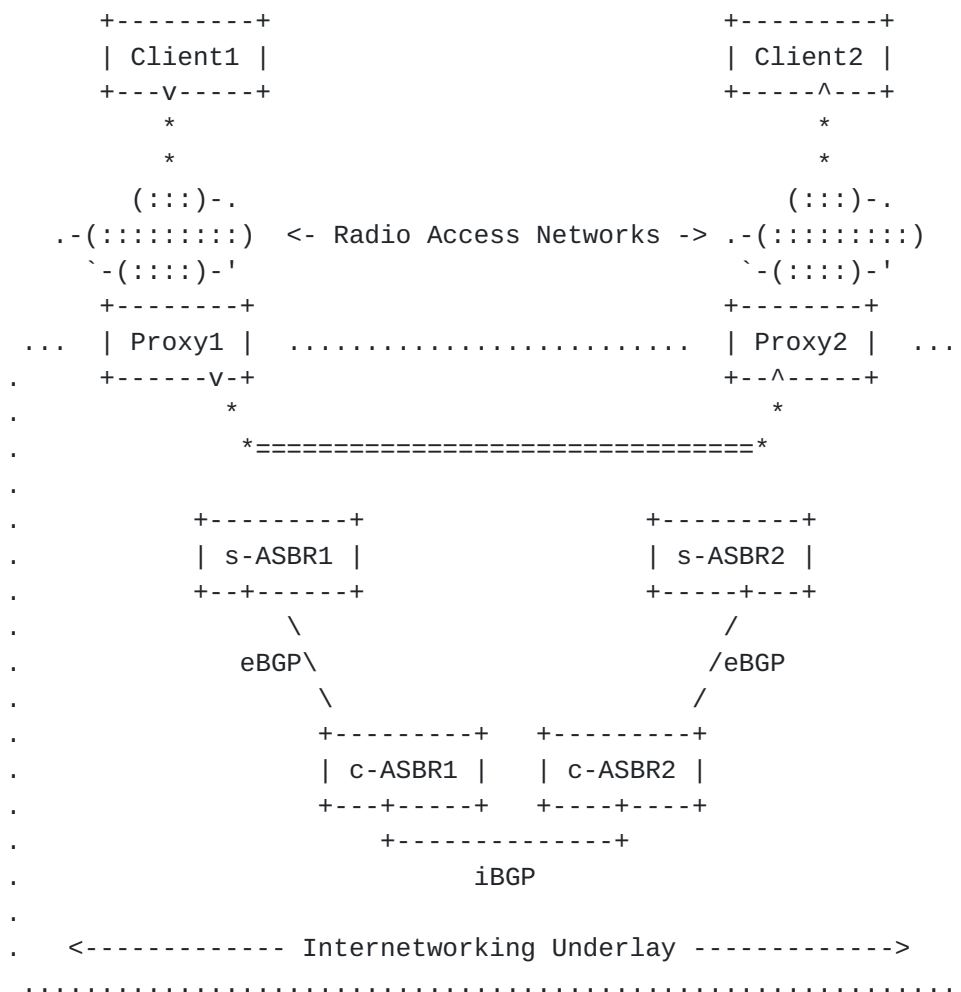
                        Figure 4: Optimized Route

6. BGP Protocol Considerations

   The number of eBGP peering sessions that each c-ASBR must service is
   proportional to the number of s-ASBRs in the system. Network
   emulations with lightweight virtual machines have shown that a single
   c-ASBR can service at least 100 eBGP peerings from s-ASBRs that each
   advertise 10K MNP routes (i.e., 1M total). It is expected that
   robust c-ASBRs can service many more peerings than this - possibly by
   multiple orders of magnitude. But even assuming a conservative
   limit, the number of s-ASBRs could be increased by also increasing
   the number of c-ASBRs. Since c-ASBRs also peer with each other using
   iBGP, however, larger-scale c-ASBR deployments may need to employ an
   adjunct facility such as BGP Route Reflectors (RRs)[RFC4456].

   The number of aircraft in operation at a given time worldwide is
   likely to be significantly less than 1M, but we will assume this
   number for a worst-case analysis. Assuming a worst-case average 1

hour flight profile from gate-to-gate with 10 service region
transitions per flight, the entire system will need to service at
most 10M BGP updates per hour (2778 updates per second). This number
is within the realm of the peak BGP update messaging seen in the
global public Internet today [BGP2]. Assuming a BGP update message
size of 100 bytes (800bits), the total amount of BGP control message
traffic to a single c-ASBR will be less than 2.5Mbps which is a
nominal rate for modern data links.

Industry standard BGP routers provide configurable parameters with
conservative default values. For example, the default hold time is
90 seconds, the default keepalive time is 1/3 of the hold time, and
the default MinRouteAdvertisementinterval is 30 seconds for eBGP
peers and 5 seconds for iBGP peers (see Section 10 of [RFC4271]).
For the simple mobile routing system described herein, these
parameters can and should be set to more aggressive values to support
faster neighbor/link failure detection and faster routing protocol
convergence times. For example, a hold time of 3 seconds and a
MinRouteAdvertisementinterval of 0 seconds for both iBGP and eBGP.

Each c-ASBR will be using eBGP both in the ATN/IPS and the
Internetworking Underlay with the ATN/IPS unicast IPv6 routes
resolving over Internetworking Underlay routes. Consequently,
c-ASBRs and potentially s-ASBRs will need to support separate local
ASes for the two BGP routing domains and routing policy or assure
routes are not propagated between the two BGP routing domains. From
a conceptual and operational standpoint, the implementation should
provide isolation between the two BGP routing domains (e.g., separate
BGP instances).

7. Implementation Status

The BGP routing topology described in this document has been modeled
in realistic network emulations showing that at least 1 million MNPs
can be propagated to each c-ASBR even on lightweight virtual
machines. No BGP routing protocol extensions need to be adopted.

8. IANA Considerations

This document does not introduce any IANA considerations.

9. Security Considerations

ATN/IPS ASBRs on the open Internet are susceptible to the same attack
profiles as for any Internet nodes. For this reason, ASBRs should
employ physical security and/or IP securing mechanisms such as IPsec
[RFC4301], TLS [RFC5246], etc.

ATN/IPS ASBRs present targets for Distributed Denial of Service
(DDoS) attacks. This concern is no different than for any node on
the open Internet, where attackers could send spoofed packets to the
node at high data rates. This can be mitigated by connecting ATN/IPS
ASBRs over dedicated links with no connections to the Internet and/or
when ASBR connections to the Internet are only permitted through
well-managed firewalls.

ATN/IPS s-ASBRs should institute rate limits to protect low data rate
aviation data links from receiving DDoS packet floods.

This document does not include any new specific requirements for
mitigation of DDoS.

10. Acknowledgements

This work is aligned with the FAA as per the SE2025 contract number
DTFAWA-15-D-00030.

This work is aligned with the NASA Safe Autonomous Systems Operation
(SASO) program under NASA contract number NNA16BD84C.

This work is aligned with the Boeing Information Technology (BIT)
MobileNet program.

11. References

11.1. Normative References

[RFC4271]   Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A
            Border Gateway Protocol 4 (BGP-4)", RFC 4271,
            DOI 10.17487/RFC4271, January 2006,
            <https://www.rfc-editor.org/info/rfc4271>.

[RFC4443]   Conta, A., Deering, S., and M. Gupta, Ed., "Internet
            Control Message Protocol (ICMPv6) for the Internet
            Protocol Version 6 (IPv6) Specification", STD 89,
            RFC 4443, DOI 10.17487/RFC4443, March 2006,
            <https://www.rfc-editor.org/info/rfc4443>.

[RFC4456]   Bates, T., Chen, E., and R. Chandra, "BGP Route
            Reflection: An Alternative to Full Mesh Internal BGP
            (IBGP)", RFC 4456, DOI 10.17487/RFC4456, April 2006,
            <https://www.rfc-editor.org/info/rfc4456>.

   [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6
              (IPv6) Specification", STD 86, RFC 8200,
              DOI 10.17487/RFC8200, July 2017,
              <https://www.rfc-editor.org/info/rfc8200>.

11.2. Informative References

   [BGP]   Huston, G., "BGP in 2015, http://potaroo.net", January
              2016.

   [BGP2]   Huston, G., "BGP Instability Report,
              http://bgpupdates.potaroo.net/instability/bgpupd.html",
              May 2017.

   [CBB]   Dul, A., "Global IP Network Mobility using Border Gateway
              Protocol (BGP), http://www.quark.net/docs/
              Global_IP_Network_Mobility_using_BGP.pdf", March 2006.

   [I-D.templin-aerolink]
              Templin, F., "Asymmetric Extended Route Optimization
              (AERO)", draft-templin-aerolink-82 (work in progress), May
              2018.

   [ICAO]   ICAO, I., "http://www.icao.int/Pages/default.aspx",
              February 2017.

   [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P.
              Traina, "Generic Routing Encapsulation (GRE)", RFC 2784,
              DOI 10.17487/RFC2784, March 2000,
              <https://www.rfc-editor.org/info/rfc2784>.

   [RFC4301] Kent, S. and K. Seo, "Security Architecture for the
              Internet Protocol", RFC 4301, DOI 10.17487/RFC4301,
              December 2005, <https://www.rfc-editor.org/info/rfc4301>.

   [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security
              (TLS) Protocol Version 1.2", RFC 5246,
              DOI 10.17487/RFC5246, August 2008,
              <https://www.rfc-editor.org/info/rfc5246>.

   [RFC6793] Vohra, Q. and E. Chen, "BGP Support for Four-Octet
              Autonomous System (AS) Number Space", RFC 6793,
              DOI 10.17487/RFC6793, December 2012,
              <https://www.rfc-editor.org/info/rfc6793>.

Appendix A.   Change Log

   << RFC Editor - remove prior to publication >>

   Changes from -07 to -08:

   o Removed suggestion to use private ASNs

   o Ran spelling checker and corrected errors

   o Re-worked Section 3 final two paragraphs on scaling

   o Stated Internetwork underlay as being out of scope for this
      document

Authors' Addresses

   Fred L. Templin (editor)
   Boeing Research & Technology
   P.O. Box 3707
   Seattle, WA 98124
   USA

   Email: fltemplin@acm.org

   Greg Saccone
   Boeing Research & Technology
   P.O. Box 3707
   Seattle, WA 98124
   USA

   Email: gregory.t.saccone@boeing.com

   Gaurav Dawra
   LinkedIn
   USA

   Email: gdawra.ietf@gmail.com

   Acee Lindem
   Cisco Systems, Inc.
   USA

   Email: acee@cisco.com

   Victor Moreno
   Cisco Systems, Inc.
   USA

   Email: vimoreno@cisco.com