

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 5 July 2022

F. L. Templin, Ed.
G. Saccone
Boeing Research & Technology
G. Dawra
LinkedIn
A. Lindem
V. Moreno
Cisco Systems, Inc.
1 January 2022

A Simple BGP-based Mobile Routing System for the Aeronautical
Telecommunications Network
draft-ietf-rtgwg-atn-bgp-12

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.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <https://datatracker.ietf.org/drafts/current/>. Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress." This Internet-Draft will expire on 5 July 2022.

Templin, et al.

Expires 5 July 2022

[Page 1]

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.
This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Terminology	6
3. ATN/IPS Routing System	8
4. ATN/IPS (Radio) Access Network (ANET) Model	12
5. ATN/IPS Route Optimization	14
6. BGP Protocol Considerations	16
7. Stub AS Mobile Routing Services	18
8. Implementation Status	18
9. IANA Considerations	18
10. Security Considerations	18
11. Acknowledgements	19
12. References	19
12.1. Normative References	19
12.2. Informative References	20
Appendix A. BGP Convergence Considerations	22
Appendix B. Change Log	22
Authors' Addresses	22

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][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 configured over one or more Internetworking underlays ("INETS") maintained by aeronautical network service providers such as ARINC, SITA and Inmarsat. The Overlay Multilink Network Interface (OMNI) [I-D.templin-6man-omni] provides a Non-Broadcast, Multiple Access (NBMA) virtual link that spans the entire ATN/IPS. Each aircraft connects to the OMNI link via an OMNI interface configured over the aircraft's underlying physical and/or virtual access network interfaces.

Each underlying INET comprises one or more "partitions" where all nodes within a partition can exchange packets with all other nodes, i.e., the partition is connected internally. There is no requirement that any two INET partitions use the same IP protocol version nor have consistent IP addressing plans in comparison with other partitions. Instead, the OMNI link sees each partition as a "segment" of a link-layer topology concatenated by a service known as the OMNI Adaptation Layer (OAL)

[I-D.templin-6man-omni][I-D.templin-6man-aero] based on IPv6 encapsulation [RFC2473].

The IPv6 addressing architecture provides different classes of addresses, including Global Unicast Addresses (GUAs), Unique Local Addresses (ULAs) and Link-Local Addresses (LLAs) [RFC4291][RFC4193]. The ATN/IPS receives an IPv6 GUA Mobility Service Prefix (MSP) from an Internet assigned numbers authority, and each aircraft will receive a Mobile Network Prefix (MNP) delegation from the MSP that accompanies the aircraft wherever it travels. ATCs and AOCs will

likewise receive MNPs, but they would typically appear in static (not mobile) deployments such as air traffic control towers, airline headquarters, etc.

The OAL uses ULAs in the source and destination addresses of IPv6 encapsulation headers. Each ULA includes an MNP in the interface identifier ("MNP-ULA") as discussed in [I-D.templin-6man-omni]. Due to MNP delegation policies and random MN mobility properties, MNP-ULAs are generally not aggregatable in the BGP routing service and are represented as many more-specific prefixes instead of a smaller number of aggregated prefixes. In addition, OMNI link service nodes configure administratively-assigned ULAs ("ADM-ULA") that are statically-assigned and derived from a shorter ADM-ULA prefix assigned to their OMNI link partition [I-D.templin-6man-omni]. Unlike MNP-ULAs, the ADM-ULAs are persistently present and unchanging in the routing system. The BGP routing services therefore perform forwarding based on these MNP-ULAs and ADM-ULAs instead of based on the GUA MNPs themselves.

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 prefixes 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 systems in underlying INETs, 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. BGP routing is based on the ULAs found in OAL headers, i.e., it provides an adaptation layer forwarding service instead of a networking layer routing service.

Commented [DT1]: A topology diagram would aid readability, I think. I.e., showing an aircraft, ATC(s), two INETs (say v4 and v6), c-ASBRs and s-ASBRs, etc. but higher level than current Figure 1.

Commented [DT2]: Mobile network? (MNP was defined above but not MN)

Commented [DT3]: Undefined term

Commented [DT4]: I find this paragraph hard to follow without a diagram.

Commented [DT5]: Two paragraphs earlier used the term "IPv6 encapsulation headers". Does this mean the source/dest addresses of an IPv6 header encapsulated in something else? Or does this mean some other header? If the former, then I find the term "OAL headers" confusing, and would simply prefer "intermediate IPv6 headers".

Commented [DT6]: As far as I know, BGP isn't data plane, it's control plane. Forwarding is a term for the data plane. So can't follow this sentence.

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 INET using industry-standard secured encapsulations (e.g., IPsec [RFC4301], Wireguard, etc.). In particular, tunneling must be used when neighboring ASBRs are separated by multiple INET hops.

The s-ASBRs engage in external BGP (eBGP) peerings with their respective c-ASBRs, and only maintain routing table entries for the MNP-ULAs currently active within the stub AS. The s-ASBRs send BGP updates for MNP-ULA 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 within the same partition using internal BGP (iBGP) peerings over which they collaboratively maintain a full routing table for all active MNP-ULAs currently in service within the partition. 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.

When there are multiple INET partitions, the c-ASBRs of each partition use eBGP to peer with the c-ASBRs of other partitions so that the full set of ULAs for all partitions are known globally among all of the c-ASBRs. Each c/s-ASBR further configures an ADM-ULA which is taken from an ADM-ULA prefix assigned to each partition, as well as static forwarding table entries for all other OMNI link partition prefixes. Both ADM-ULAs and MNP-ULAs are used by the OAL for nested encapsulation where the inner IPv6 packet is encapsulated in an IPv6 OAL header with ULA source and destination addresses, which is then encapsulated in an IP header specific to the INET partition.

With these intra- and inter-INET BGP peerings in place, a forwarding plane spanning tree is established that properly covers the entire operating domain. All nodes in the network can be visited using strict spanning tree hops, but in many instances this may result in longer paths than are necessary. The AERO and OMNI services [I-D.templin-6man-aero][I-D.templin-6man-omni] provide mechanisms for discovering and utilizing (route-optimized) shortcuts that do not always follow strict spanning tree paths.

Commented [DT7]: This sentence is helpful and I would like to see it moved earlier, to the first full paragraph at top of page 4 (where I had comments about it being hard to understand).

Commented [DT8]: This wording implies a normative dependency, but they're informative references. If they're just examples, then the wording here should be updated. If they're normative, then does the RTGWG want to take a normative dependency on (currently) non-IETF drafts?

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 one or more Internetworking underlays.

Internetworking underlay ("INET")

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

Example INET service providers for civil aviation include ARINC, SITA and Inmarsat.

(Radio) Access Network ("ANET")

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 outside INETs. The term ANET is intended in the same spirit as for radio-based Internet service provider networks (e.g., cellular operators), but can also refer to ground-domain networks that connect AOCs and ATCs.

partition (or "segment")

A fully-connected internal subnetwork of an INET in which all nodes can communicate with all other nodes within the same partition using the same IP protocol version and addressing plan. Each INET consists of one or more partitions.

Overlay Multilink Network Interface (OMNI)

A virtual layer 2 bridging service that presents an ATN/IPS overlay unified link view even though the underlay may consist of multiple INET partitions. The OMNI virtual link is manifested through nested encapsulation in which GUA-addressed IPv6 packets from the ATN/IPS are first encapsulated in ULA-addressed IPv6 headers which are then forwarded to the next hop using INET encapsulation if necessary. Forwarding over the OMNI virtual link is therefore based on ULAs instead of GUAs. In this way, packets sent from a source can be conveyed over the OMNI virtual link even though there may be many underlying INET partitions in the path to the destination.

OMNI Adaptation Layer (OAL)

A middle layer below the IPv6 layer but above the INET layer that applies IPv6-in-IPv6 encapsulation prior to INET encapsulation. The IPv6 encapsulation header inserted by the OAL uses ULAs instead of GUAs. Further details on OMNI and the OAL are found in [I-D.templin-6man-omni].

OAL Autonomous System

A "hub-of-hubs" autonomous system maintained through peerings between the core autonomous systems of different OMNI virtual link partitions.

Core Autonomous System Border Router (c-ASBR)

A BGP router located in the hub of the INET partition hub-and-spokes overlay network topology.

Core Autonomous System

The "hub" autonomous system maintained by all c-ASBRs within the same partition.

Stub Autonomous System Border Router (s-ASBR)

A BGP router configured as a spoke in the INET partition hub-and-spokes overlay network topology.

Stub Autonomous System

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

Commented [DT9]: Either here or in a diagram like the one I asked for above, it would helpful to note what roles are the endpoints of the OAL. E.g., does it go between Clients and ASBR's? Or between a Proxy/Server and ASBR's?

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 or fixed network.

Proxy/Server

An ANET/INET border node that acts as a transparent intermediary between Clients and s-ASBRs. From the Client's perspective, the Proxy/Server presents the appearance that the Client is communicating directly with the s-ASBR. From the s-ASBR's perspective, the Proxy/Server 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 ATN/IPS routing system comprises a private BGP instance coordinated in an overlay network via tunnels between neighboring ASBRs over one or more underlying INETs. The overlay does not interact with the underlying INET BGP routing systems, and only a small and unchanging set of MSPs are advertised externally instead of the full dynamically changing set of MNPs.

Within each INET partition, each s-ASBR^s connects a stub AS to the INET partition core using a distinct 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 INET partition core AS, and use a shared core ASN. Unique ASNs are assigned according to the standard 32-bit ASN format [RFC4271][RFC6793]. Since the BGP instance does not connect with any INET BGP routing systems, the ASNs assigned **need not be coordinated with IANA** and can in fact coincide with values that are assigned in other domains. The only requirement is that ASNs must not be duplicated within the ATN/IPS routing system itself.

The c-ASBRs use iBGP to maintain a synchronized consistent view of all active MNP-ULAs currently in service within the INET partition. Figure 1 below represents the reference INET partition 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

Commented [DT10]: Remove "s"?

Commented [DT11]: Who coordinates them then, to ensure that are unique within the ATN/IPS routing system?

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.

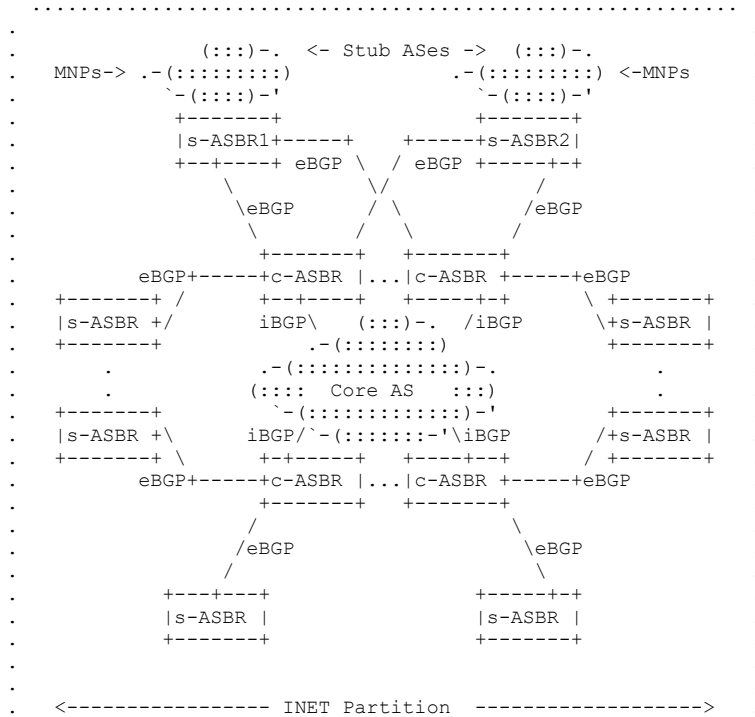


Figure 1: INET Partition Reference Deployment

In the reference deployment, each s-ASBR maintains routes for active MNP-ULAs that currently belong to its stub AS. In response to "Inter-domain" mobility events, each s-ASBR ~~will~~ dynamically announces new MNP-ULAs and withdraws departed MNP-ULAs 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 MNP-ULAs that are currently active, and packets destined to all other MNP-ULAs 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 MNP-ULAs 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 MNP-ULAs currently within their stub ASes) and they forward all other packets to c-ASBRs which have full topology knowledge. Each s-ASBR and c-ASBR configures an ADM-ULA that is aggregatable within an INET partition, and each partition configures a unique ADM-ULA prefix that is permanently announced into the routing system. The core ASes of each INET partition are joined together through external BGP peerings. The c-ASBRs of each partition establish external peerings with the c-ASBRs of other partitions to form a "core-of-cores" OMNI link AS. The OMNI link AS contains the global knowledge of all MNP-ULAs deployed worldwide, and supports ATN/IPS overlay communications between nodes located in different INET partitions by virtue of OAL encapsulation. OMNI link nodes can then navigate to ASBRs by including an ADM-ULA or directly to an end system by including an MNP-ULA in the destination address of an OAL-encapsulated packet (see: [I-D.templin-6man-aero]). Figure 2 shows a reference OAL topology.

Commented [DT12]: Grammar: delete word

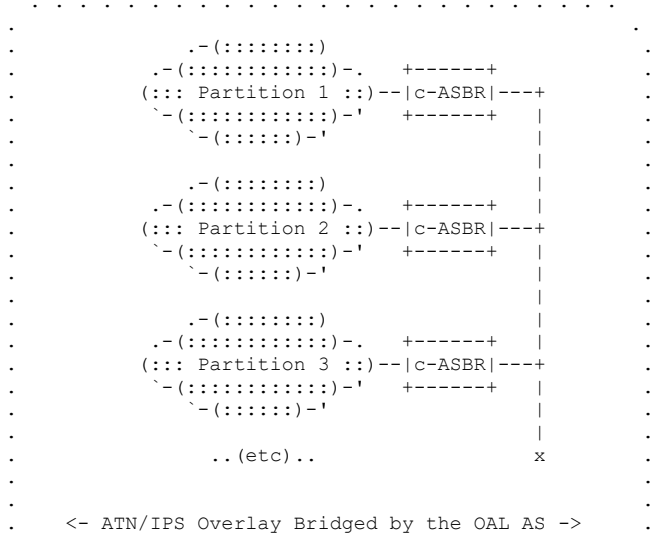


Figure 2: Spanning Partitions with the OAL

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 MNP-ULAs 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, with potential for supporting many millions of mobile networks or more. In this way, each set of c-ASBRs services a specific set of MSPs, 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. For larger-scale applications (such as unmanned air vehicles and terrestrial vehicles) even larger scales can be accommodated by adding more c-ASBRs.

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

(Radio) Access Networks (ANETs) connect end system Clients such as aircraft, ATCs, AOCs etc. to the ATN/IPS routing system. Clients may connect to multiple ANETs at once, for example, when they have both satellite and cellular data links activated simultaneously. Clients configure an Overlay Multilink Network (OMNI) Interface [I-D.templin-6man-omni] over their underlying ANET interfaces as a connection to an NBMA virtual link (manifested by the OAL) that spans the entire ATN/IPS. Clients may further move between ANETs in a manner that is perceived as a network layer mobility event. Clients could therefore employ a multilink/mobility routing service such as those discussed in Section 7.

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/Server at the ANET/INET boundary.

Figure 3 shows the ATN/IPS ANET model where Clients connect to ANETs via aviation data links. Clients register their ANET addresses with a nearby s-ASBR, where the registration process may be brokered by a Proxy/Server at the edge of the ANET.

Commented [DT13]: As worded, this sounds like a normative reference.

Commented [DT14]: Related to my questions about whether omni draft references are normative or merely examples, if an example then could PMIP (RFC 5213) be used as another example?

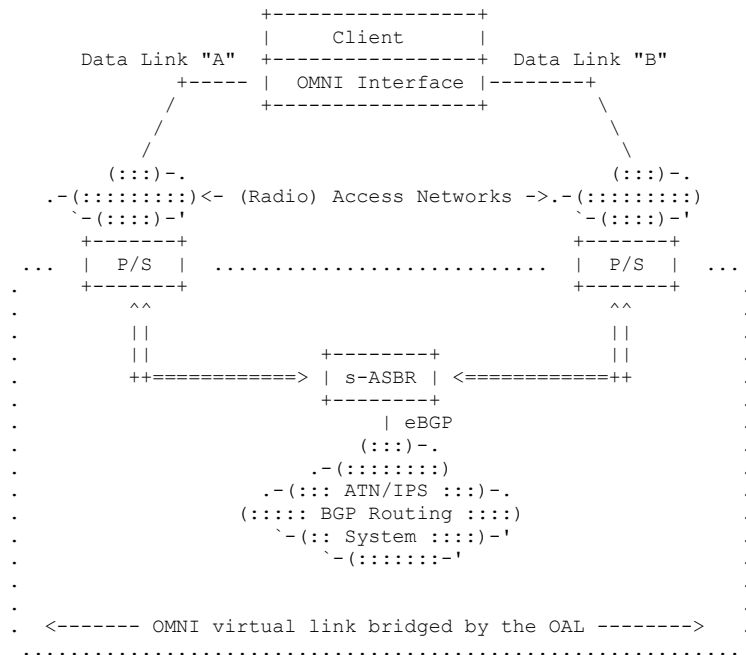


Figure 3: ATN/IPS ANET Architecture

When a Client logs into an ANET it specifies a nearby s-ASBR that it has selected to connect to the ATN/IPS. The login process is transparently brokered by a Proxy/Server at the border of the ANET which then conveys the connection request to the s-ASBR via tunneling across the OMNI virtual link. Each ANET border Proxy/Server is also equally capable of serving in the s-ASBR role so that a first on-link Proxy/Server can be selected as the s-ASBR while all others perform the Proxy/Server role in a hub-and-spokes arrangement. An on-link Proxy/Server is selected to serve the s-ASBR role when it receives a control message from a Client requesting that service. The Client can coordinate with a network-based s-ASBR over additional ANETs after it has already coordinated with a first-hop Proxy/Server over a first ANET. Selection of a network-based s-ASBR could be through an address discovered through a first ANET Proxy/Server, through consulting a geographically-keyed static host file, through a

Templin, et al. Expires 5 July 2022 [Page 13]

Commented [DT15]: "connects to"? Logs into sounds like a human does it manually.

Commented [DT16]: Is the choice implementation specific? Or does it need to be interoperable? The phrase "discovered through a Proxy/Server" sounds like you need a protocol.

DNS lookup, through a network query response, etc. The s-ASBR then registers the addresses of the additional ANET Proxy/Server as the address for the Client over each distinct Client interface. If the Client connects to multiple ANETs, the s-ASBR will register the addresses of all Proxy/Servers as addresses through which the Client can be reached.

The s-ASBR represents all of its active Clients as MNP-ULA 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 MNP-ULAs of its active Clients and withdraws the MNP-ULAs of its departed Clients via BGP updates to c-ASBRs, which further propagate the MNP-ULAs to other c-ASBRs within the OAL AS. Since Clients are expected to remain associated with their current s-ASBR for extended periods, the level of MNP-ULA 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 spanning tree segments in the forwarding path. In many cases, it would be desirable to eliminate extraneous spanning tree segments from this "dogleg" route so that packets can traverse a minimum number of tunneling hops across the OMNI virtual link. ATN/IPS end systems could therefore employ a route optimization service according to the mobility service employed (see: Section 7).

A route optimization example is shown in Figure 4 and Figure 5 below. In the first figure, multiple spanning tree 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.

Commented [DT17]: Registers to what, using what protocol?

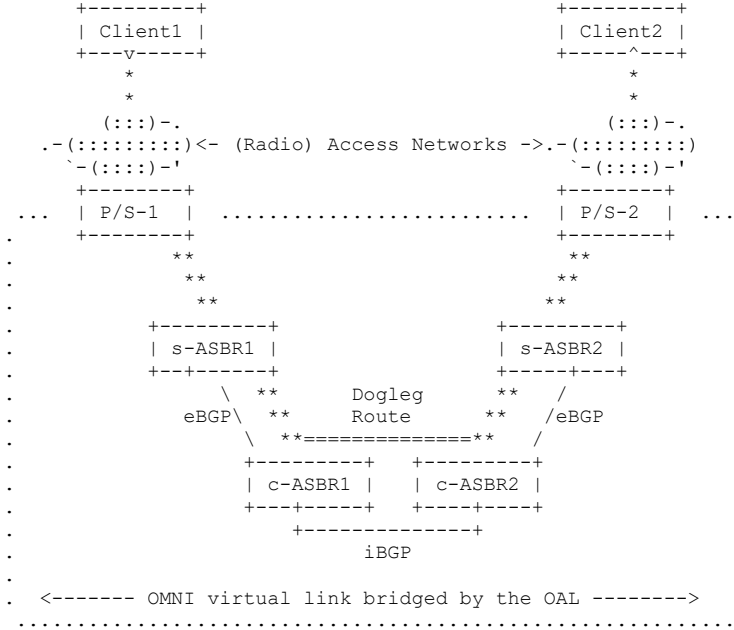



Figure 4: Dogleg Route Before Optimization
Expires 5 July 2022



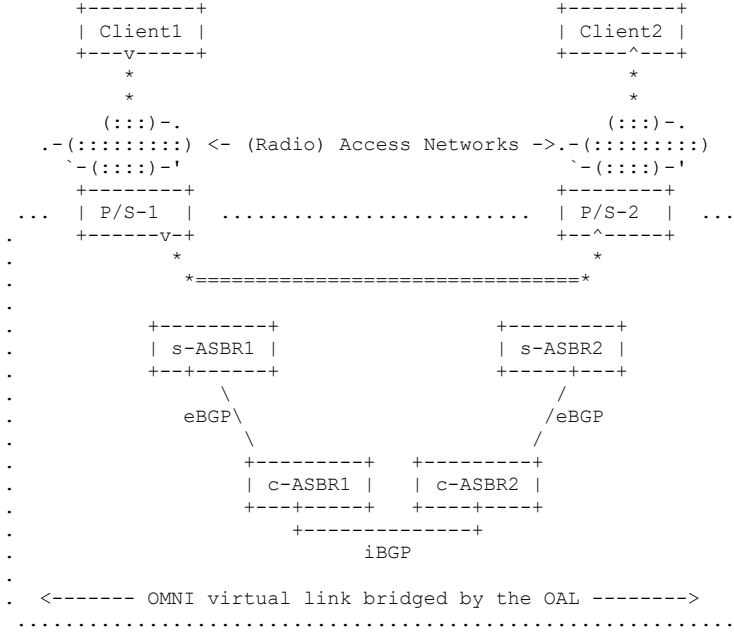


Figure 5: Optimized Route

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 its local partition. 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-ULA 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 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. Instead of adjusting BGP default time values, BGP routers can use the Bidirectional Forwarding Detection (BFD) protocol [RFC5880] to quickly detect link failures that don't result in interface state changes, BGP peer failures, and administrative state changes. BFD is important in environments where rapid response to failures is required for routing reconvergence and, hence, communications continuity.

Each c-ASBR will be using eBGP both in the ATN/IPS and the INET with the ATN/IPS unicast IPv6 routes resolving over INET 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).

ADM-ULAs and MNP-ULAs begin with fd00::/8 followed by a pseudo-random 40-bit global ID to form the prefix [ULA]::/48, along with a 16-bit OMNI link identifier '*' to form the prefix [ULA*]::/64. Each individual address taken from [ULA*]::/64 includes additional routing information in the interface identifier. For example, for the MNP 2001:db8:1:0::/56, the resulting MNP-ULA is [ULA*]:2001:db8:1:0/120, and for the administrative address 1001:2002/16 the ADM-ULA is [ULA*]::1001:2002/112 (see: [I-D.templin-6man-omni] for further

Commented [DT18]: But also from a correctness standpoint right? (since AS numbers can be reused and there may be collisions between the two)

details). This gives rise to a BGP routing system that must accommodate large numbers of long and non-aggregatable MNP-ULA prefixes as well as moderate numbers of long and semi-aggregatable ADM-ULA prefixes. The system is kept stable and scalable through the s-ASBR / c-ASBR hub-and-spokes topology which ensures that mobility-related churn is not exposed to the core.

7. Stub AS Mobile Routing Services

Stub ASes maintain intradomain routing information for mobile node clients, and are responsible for all localized mobility signaling without disturbing the BGP routing system. Clients can enlist the services of a candidate mobility service such as Mobile IPv6 (MIPv6) [RFC6275], LISP [I-D.ietf-lisp-rfc6830bis] and AERO [I-D.templin-6man-aero] according to the service offered by the stub AS. Further details of mobile routing services are out of scope for this document.

8. Implementation Status

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

9. IANA Considerations

This document does not introduce any IANA considerations.

10. 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], WireGuard, 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.

Commented [DT19]: Grammar: "or"?

BGP protocol message exchanges and control message exchanges used for route optimization must be secured to ensure the integrity of the system-wide routing information base.

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

11. 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 Commercial Airplanes (BCA) Internet of Things (IoT) and autonomy programs.

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

The following individuals contributed insights that have improved the document: Ahmad Amin, Erik Kline, Hubert Kuenig, Tony Li, Alexandre Petrescu, Pascal Thubert, Tony Whyman.

12. References

12.1. Normative References

- [RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, DOI 10.17487/RFC2473, December 1998, <<https://www.rfc-editor.org/info/rfc2473>>.
- [RFC4193] Hinden, R. and B. Haberman, "Unique Local IPv6 Unicast Addresses", RFC 4193, DOI 10.17487/RFC4193, October 2005, <<https://www.rfc-editor.org/info/rfc4193>>.
- [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>>.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", RFC 4291, DOI 10.17487/RFC4291, February 2006, <<https://www.rfc-editor.org/info/rfc4291>>.

- [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>>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <<https://www.rfc-editor.org/info/rfc5880>>.
- [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>>.

12.2. Informative References

- [ATN] Maiolla, V., "The OMNI Interface - An IPv6 Air/Ground Interface for Civil Aviation, IETF Liaison Statement #1676, <https://datatracker.ietf.org/liaison/1676/>", 3 March 2020.
 - [ATN-IPS] WG-I, ICAO., "ICAO Document 9896 (Manual on the Aeronautical Telecommunication Network (ATN) using Internet Protocol Suite (IPS) Standards and Protocol), Draft Edition 3 (work-in-progress)", 10 December 2020.
 - [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.
- Templin, et al. Expires 5 July 2022 [Page 20]

[I-D.ietf-lisp-rtc6830bis]

Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos, "The Locator/ID Separation Protocol (LISP)", Work in Progress, Internet-Draft, draft-ietf-lisp-rtc6830bis-36, 18 November 2020, <<https://www.ietf.org/archive/id/draft-ietf-lisp-rtc6830bis-36.txt>>.

[I-D.templin-6man-aero]

Templin, F. L., "Automatic Extended Route Optimization (AERO)", Work in Progress, Internet-Draft, draft-templin-6man-aero-37, 15 November 2021, <<https://www.ietf.org/archive/id/draft-templin-6man-aero-37.txt>>.

[I-D.templin-6man-omni]

Templin, F. L. and T. Whyman, "Transmission of IP Packets over Overlay Multilink Network (OMNI) Interfaces", Work in Progress, Internet-Draft, draft-templin-6man-omni-51, 15 November 2021, <<https://www.ietf.org/archive/id/draft-templin-6man-omni-51.txt>>.

[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>>.

[RFC6275] Perkins, C., Ed., Johnson, D., and J. Arkko, "Mobility Support in IPv6", RFC 6275, DOI 10.17487/RFC6275, July 2011, <<https://www.rfc-editor.org/info/rfc6275>>.

[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>>.

Experimental evidence has shown that BGP convergence time required ~~for~~ when an MNP-ULA is asserted at a new location or withdrawn from an old location can be several hundred milliseconds even under optimal AS peering arrangements. This means that packets in flight destined to an MNP-ULA route that has recently been changed can be (mis)delivered to an old s-ASBR after a Client has moved to a new s-ASBR.

To address this issue, the old s-ASBR can maintain temporary state for a "departed" Client that includes an OAL address for the new s-ASBR. The OAL address never changes since ASBRs are fixed infrastructure elements that never move. Hence, packets arriving at the old s-ASBR can be forwarded to the new s-ASBR while the BGP routing system is still undergoing reconvergence. Therefore, as long as the Client associates with the new s-ASBR before it departs from the old s-ASBR (while informing the old s-ASBR of its new location) packets in flight during the BGP reconvergence window are accommodated without loss.

Appendix B. Change Log

<< RFC Editor - remove prior to publication >>
Differences from earlier versions:
* Submit for RFC publication.

Authors' Addresses

Fred L. Templin (editor)
Boeing Research & Technology
P.O. Box 3707
Seattle, WA 98124
United States of America
Email: fltemplin@acm.org
Greg Saccone
Boeing Research & Technology
P.O. Box 3707
Seattle, WA 98124
United States of America
Email: gregory.t.saccone@boeing.com

Templin, et al.

Expires 5 July 2022

[Page 22]

Commented [DT20]: grammar

Internet-Draft

BGP for ATN/IPS

January 2022

Gaurav Dawra

LinkedIn

United States of America

Email: gdawra.ietf@gmail.com

Acee Lindem

Cisco Systems, Inc.

United States of America

Email: acee@cisco.com

Victor Moreno

Cisco Systems, Inc.

United States of America

Email: vimoreno@cisco.com

Templin, et al.

Expires 5 July 2022

[Page 23]

