

Network Working Group
Internet-Draft
Intended status: Informational
Expires: 18 September 2025

F. L. Templin, Ed.
G. Saccone
Boeing Research & Technology
G. Dawra
LinkedIn
A. Lindem
LabN Consulting LLC
V. Moreno
Cisco Systems, Inc.
17 March 2025

A Simple BGP-based Mobile Routing System for the Aeronautical
Telecommunications Network
draft-ietf-rtgwg-atn-bgp-28

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 IP-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 18 September 2025.

Copyright Notice

Copyright (c) 2025 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	7
3. ATN/IPS Routing System	10
4. ATN/IPS (Radio) Access Network (ANET) Model	14
5. ATN/IPS Route Optimization	16
6. BGP Protocol Considerations	18
7. Stub AS Mobile Routing Services	20
8. Implementation Status	20
9. IANA Considerations	20
10. Security Considerations	20
10.1. Public Key Infrastructure (PKI) Considerations	21
11. Acknowledgements	21
12. References	22
12.1. Normative References	22
12.2. Informative References	23
Appendix A. BGP Convergence Considerations	25
Appendix B. AERO/OMNI Spanning Tree Properties	25
Appendix C. Change Log	26
Authors' Addresses	26

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 some worldwide ATM domains.

The International Civil Aviation Organization (ICAO) is now engaged in 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 Internetworking service supporting pervasive ATM for Air Traffic Controllers (ATC), Airline Operations Controllers (AOC), and all commercial aircraft worldwide. The ATN/IPS will require a new mobile routing service as a core element of the architecture. 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. Furthermore, continuous performance-intensive control messaging services such as BGP peering sessions must be carried only over the well-connected secured ground domain ATN/IPS network and never over low-end aviation data links.

The ATN/IPS is an IP-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-omni3] uses an adaptation layer encapsulation to create a Non-Broadcast, Multiple Access (NBMA) virtual link spanning 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 each INET partition uses the same IP protocol version nor has consistent IP addressing plans in relation to other partitions. Instead, the OMNI link sees each partition as a "segment" of a lower layer topology concatenated by a service known as the OMNI Adaptation Layer (OAL) [I-D.templin-6man-omni3] 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]. A new class known as Multilink-Local Addresses (MLAs) [I-D.templin-6man-mla] is also necessary to support the OMNI virtual link abstraction over an end system's multiple underlying interfaces.

The ATN/IPS receives IPv6 GUA Mobility Service Prefixes (MSPs) from an Internet assigned numbers authority, and each aircraft will receive a Provider-Independent (PI) Mobile Network Prefix (MNP) delegation from an MSP that travels with the aircraft. ATCs and AOCs will instead configure Provider-Aggregated (PA) addresses from Stable Network Prefixes SNPs also taken from an MSP or Foreign Network Prefix (FNP) exposed by different networks. Services that assign addresses taken from SNPs typically appear in static (not mobile) deployments such as air traffic control towers, airline headquarters, etc., while FNP correspondents could be in either static or mobile environments in foreign networks. (Note that while IPv6 GUAs are assumed for ATN/IPS, IPv4 with public/private addresses could also be used.)

The ATN/IPS service uses ULAs for local network addressing. Each ULA includes a prefix beginning with "fd00::/8" followed by a 40-bit Global ID and a 16-bit Subnet ID as "fd{Global ID}:{Subnet ID}::/64". Each aircraft Client assigns the Subnet Router Anycast (SRA) GUA for each of its delegated MNPs and associates them with SNP GUAs delegated by First-Hop Segment (FHS) Proxy/Servers on their underlying interfaces while coordinating with a Mobility Anchor Point (MAP) Proxy/Server to register MNP SRA GUAs in the routing system. Due to MNP delegation policies and random node mobility properties, MNPs are generally not aggregable in the BGP routing service and are often represented as many more-specific prefixes instead of a smaller number of aggregated prefixes.

BGP routing service infrastructure nodes additionally configure SNP SRA GUAs that are statically-assigned and derived from a shorter SNP assigned to their BGP network partitions. Unlike MNPs, the SNP SRA GUAs are persistently present and unchanging in the routing system. The BGP routing services therefore establish forwarding table entries based on these MNPs and SNPs which support adaptation layer (not network layer) routing and addressing.

When a source Client's OMNI interface forwards an original IP packet, the OAL performs IPv6 encapsulation using the source and target Client MLAs as Source and Destination Addresses and with an OMNI Routing Header (ORH) which includes SNP GUAs for the interdomain routing region. The OAL source next subjects each resulting "OAL packet" to IPv6 fragmentation and header compression, then

encapsulates each fragment in INET headers specific to the underlying Internetwork. These resulting "carrier packets" then traverse a series of Internetworks connected by OAL intermediate nodes which re-encapsulate them in new INET headers for traversal of the next Internetwork in succession. (The carrier packets may themselves be subject to IP layer fragmentation/reassembly during each successive Internetwork traversal at a layer below the OAL.) The carrier packets finally arrive at the OAL destination which performs adaptation layer decapsulation and reassembly to restore the original IP packet. A high-level ATN/IPS network diagram is shown in Figure 1:

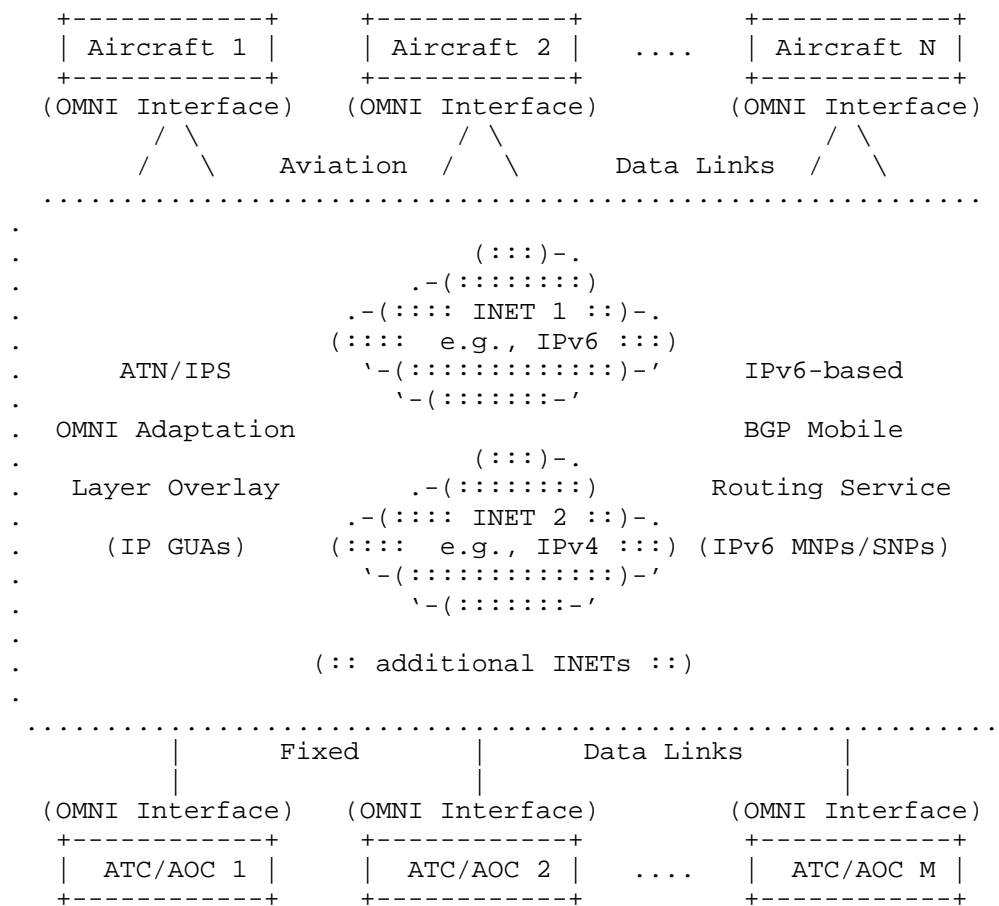


Figure 1: ATN/IPS Network Diagram

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 which must be propagated to all BGP routers in the Internet with full routing tables.

We therefore consider a routing system using an overlay network that maintains a private BGP routing protocol instance 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, and BGP routing is based on adaptation layer MNP/SNP SRA GUAs instead of network or data link layer public/private IP prefixes. This allows ASBRs to perform adaptation layer forwarding based on intermediate layer IPv6 header information instead of network layer forwarding based on upper layer IP header information or link layer forwarding based on lower layer IP header information.

The s-ASBRs for each stub AS connect to a small number of c-ASBRs via secured direct point-to-point links and/or secured tunnels (e.g., IPsec [RFC4301], etc.) over the underlying INET. Neighboring c-ASBRs should also use IP layer or lower-layer security services over their connecting links to ensure INET layer security.

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 within the same partition using internal BGP (iBGP) peerings over which they collaboratively maintain a full routing table for all active MNPs 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 SNP SRA GUA taken from a SNP assigned to each partition, as well as static forwarding table entries for all other OMNI link partition prefixes.

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. AERO [I-D.templin-6man-aero3] provides an example service for discovering and utilizing (route-optimized) shortcuts that do not always follow strict spanning tree paths.

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 term Internet Protocol (IP) refers generically to either protocol version unless specifically qualified as IPv4 [RFC0791] or IPv6 [RFC8200].

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 provide an IP-based overlay network service that connects access networks via encapsulation and forwarding over one or more Internetworking underlays.

Internetworking underlay ("INET")

A wide-area network that supports overlay network encapsulation/forwarding 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 original IP packets from ATN/IPS Clients are first encapsulated in adaptation layer 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 the adaptation layer addresses instead of the original IP addresses. 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 IP layer but above the INET layer that forwards original IP packets by applying IPv6 encapsulation, fragmentation and header compression followed by INET encapsulation to form "carrier packets". End systems that

configure OMNI interfaces act as the OAL source and destination, while intermediate systems with OMNI interfaces act as OAL forwarding nodes. Regardless of the number of intermediate systems in the path, the network layer IP TTL/Hop Limit is not decremented during (OAL layer) forwarding. Further details on OMNI and the OAL are found in [I-D.templin-6man-omni3].

OAL Autonomous System (OAL AS)

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 (Core AS)

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 (Stub AS)

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 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 IP prefix that is delegated to an ATN/IPS end system, including mobile aircraft and in some cases also ATCs, AOCs, etc. in fixed configurations. Each end system configures and responds to a Subnet Router Anycast (SRA) GUA taken from the MNP.

Stable Network Prefix (SNP)

An IPv6 prefix assigned to a c/s-ASBR in a stable configuration in the routing system. Each c/s-ASBR assigns the Subnet Router Anycast (SRA) taken from its assigned SNP, and each s-ASBR delegates GUA/ULA address pairs to aircraft Clients which associate the addresses with their delegated MNPs.

Foreign Network Prefix (FNP)

An IP prefix assigned to another link or network outside of the OMNI link domain. Relay s-ASBRs advertise the FNPs into the BGP routing service and act as routers to forward packets between MNP/SNP correspondents on the OMNI link and FNP correspondents on other links.

Mobility Service Prefix (MSP)

An aggregated IP prefix assigned to the ATN/IPS by an Internet assigned numbers authority, and from which all MNPs/SNPs are delegated (e.g., up to 2^{32} IPv6 /56 MNPs/SNPs 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 ATN/IPS routing system interacts with underlying INET BGP routing systems only through the static advertisement of a small and unchanging set of MSPs instead of the full dynamically changing set of MNPs.

Within each INET partition, each s-ASBR 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 can be assigned from the [RFC6996] 32-bit ASN space which reserves 94,967,295 numbers for private use. The ASNs must be allocated and managed by an ATN/IPS assigned numbers authority established by ICAO, which must ensure that ASNs are responsibly distributed without duplication and/or overlap.

The c-ASBRs use iBGP to maintain a synchronized consistent view of all active MNPs/SNPs currently in service within the INET partition. Figure 2 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 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.

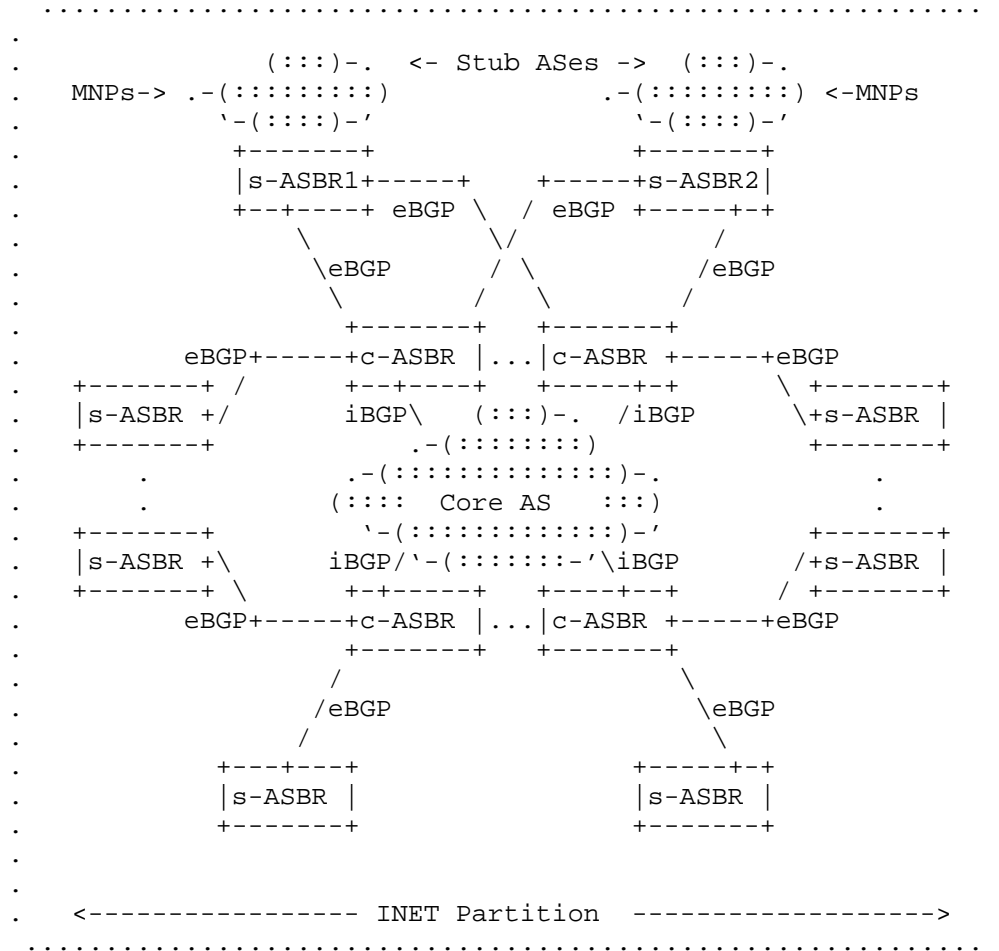


Figure 2: INET Partition Reference Deployment

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 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 maintains forwarding table entries only for the MNPs/SNPs that are currently active. If an arriving packet matches a black-hole route without matching an MNP/SNP, the c-ASBR should drop the packet and may also generate an ICMPv6 Destination Unreachable message [RFC4443], i.e., without forwarding the packet outside of the ATN/IPS overlay based on a less-specific route.

The c-ASBRs do not send BGP updates for MNPs/SNPs 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.

Each s-ASBR and c-ASBR configures an SNP SRA GUA that is aggregable within an INET partition, and each partition configures a unique SNP 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 MNPs/SNPs 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 SNP GUA or directly to an end system by including an MNP GUA in the destination address of an OAL-encapsulated packet (see: [I-D.templin-6man-aero3]). Figure 3 shows a reference OAL topology.

Figure 3: Spanning Partitions with the OAL

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, with potential for supporting many millions of mobile networks or more.

In this design, 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 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, additional c-ASBRs can be added later 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.

Consider now that the c-ASBRs provide adaptation layer gateways between independent Internetworks to form a true network-of-networks supporting the ATN/IPS overlay. This same arrangement was first envisioned by the "Catenet Model for Internetworking" [IEN48][IEN48-2] circa 1978.

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 satellite, cellular, WiFi and/or other data links activated simultaneously. Each Client configures an OMNI interface [I-D.templin-6man-omni3] over its 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 should therefore employ a multilink/mobility routing service such as those discussed in Section 7.

Clients and s-ASBRs assign MLAs to their OMNI interfaces for use as adaptation layer Source and Destination Addresses during encapsulation. Clients and s-ASBRs use the MLAs to establish Neighbor Cache Entries (NCEs) via the OMNI interface as well as for (multihop) forwarding within the Access Networks. The s-ASBRs in turn rewrite the MLAs to SNP SRA GUAs to support forwarding over the ATN/IPS interdomain region. Clients register 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 4 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.

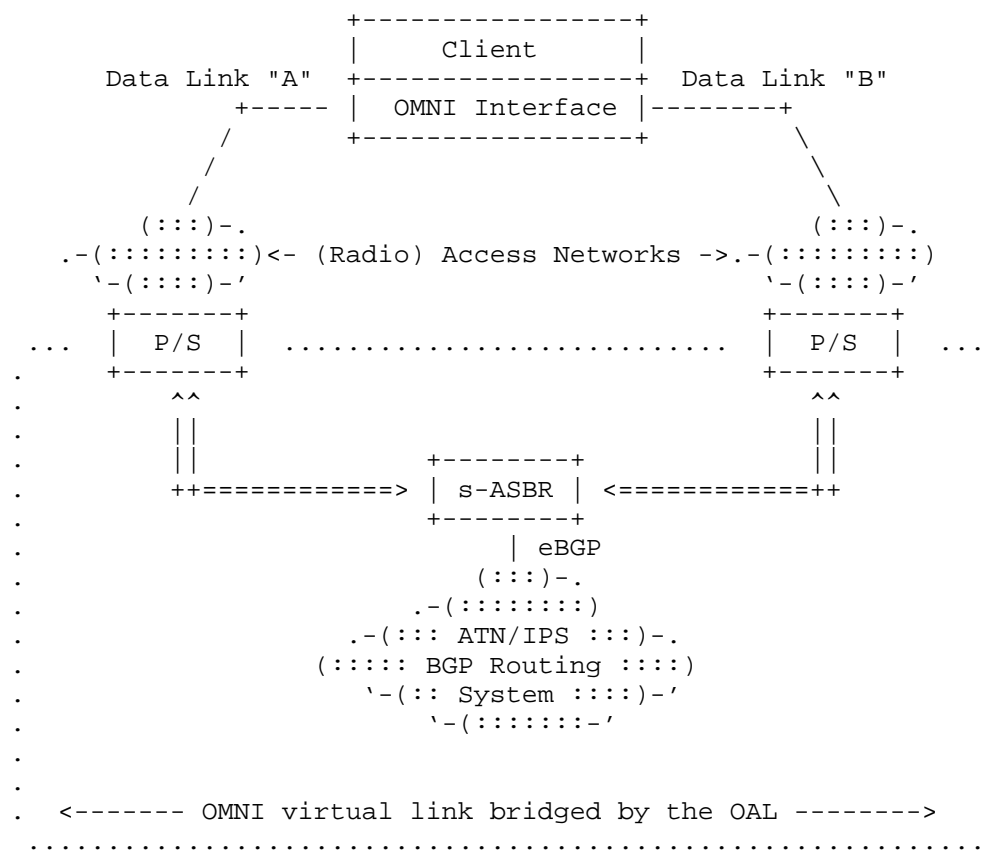


Figure 4: ATN/IPS ANET Architecture

When a Client connects to 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 adaptation layer encapsulation and forwarding 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 specifically 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. If the Client connects to multiple ANETs, the s-ASBR will register the individual ANET Proxy/Servers as conduits through which the Client can be reached. The Client then sees the s-ASBR as the "hub" in a "hub-and-spokes" arrangement with the first-hop Proxy/Servers as spokes. Selection of a network-based s-ASBR is through the discovery methods specified in relevant mobility and virtual link coordination specifications (e.g., see AERO [I-D.templin-6man-aero3] and OMNI [I-D.templin-6man-omni3]).

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 is therefore used only to advertise the set of MNPs of all its active Clients to its BGP peer c-ASBRs and not to peer with other s-ASBRs (i.e., the stub AS is a logical construct and not a physical one). The s-ASBR injects the MNPs of its active Clients and withdraws the MNPs of its departed Clients via BGP updates to c-ASBRs, which further propagate the MNPs 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 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 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 extraneous hops and/or spanning tree segments in the forwarding path. In many cases, it would be desirable to establish a "short cut" around this "dogleg" route so that packets can traverse a minimum number of forwarding 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).

Each s-ASBR provides designated routing services for only a subset of all active Clients, and instead acts as a simple Proxy/Server for other Clients. As a designated router, the s-ASBR advertises the MNPs of each of its active Clients into the ATN/IPS routing system and provides basic (unoptimized) forwarding services when necessary.

These route optimized paths are established through secured control plane messaging (i.e., over secured tunnels and/or using higher-layer control message authentications) but do not provide lower-layer security for the data plane. Data communications over these route optimized paths should therefore employ higher-layer security.

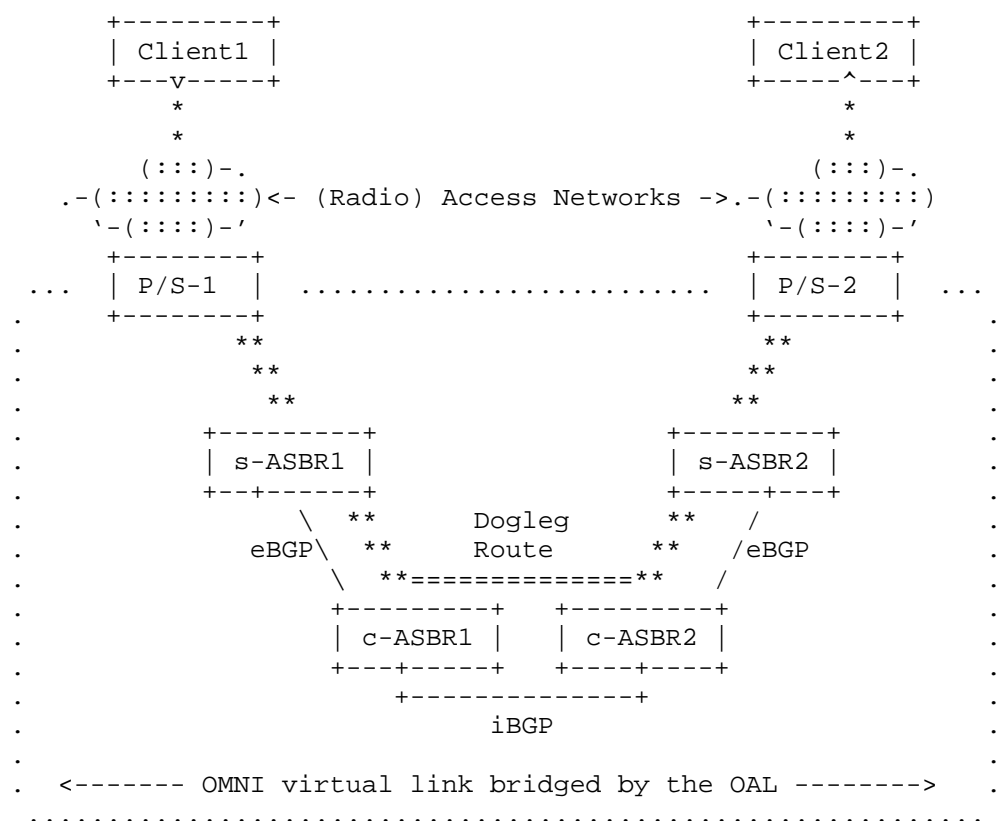


Figure 5: Dogleg Route Before Optimization

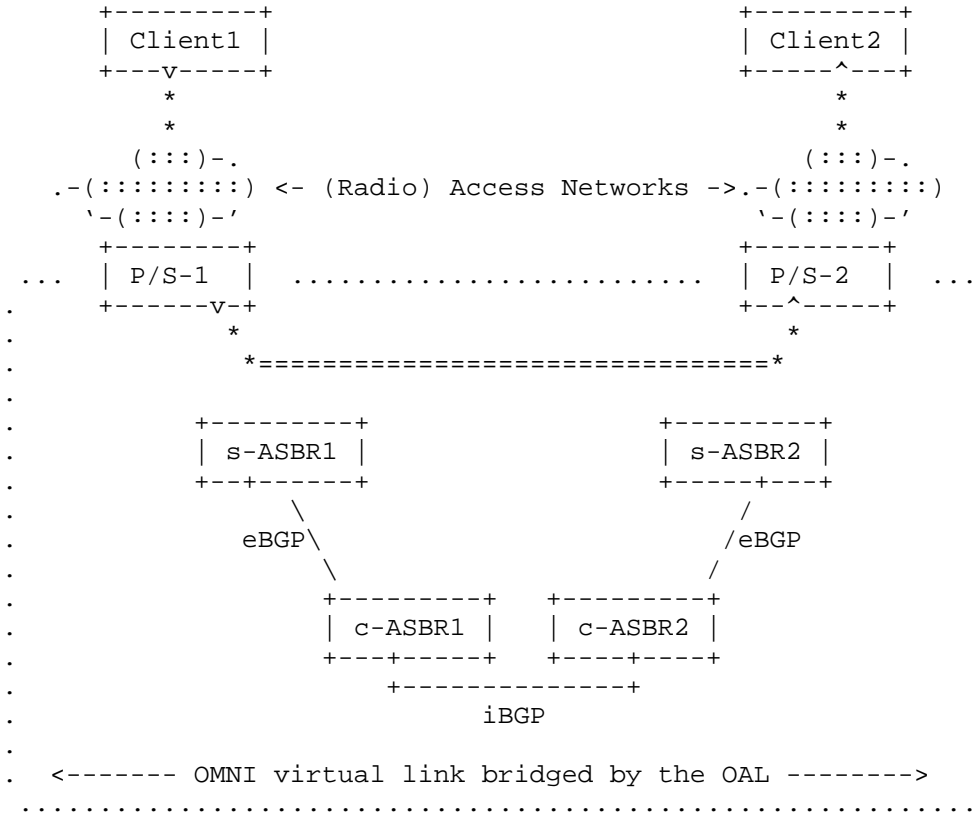


Figure 6: Optimized Route

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 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 octets (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 do not 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 IP 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, operational and correctness standpoint, the implementation should provide isolation between the two BGP routing domains (e.g., separate BGP instances).

This gives rise to a BGP routing system that must accommodate large numbers of long and non-aggregable MNP prefixes as well as moderate numbers of shorter and semi-aggregable SNP 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] or AERO [I-D.templin-6man-aero3] 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 MNPs 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], 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.

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. Security is based on IP layer security associations between peers which ensure confidentiality, integrity and authentication over secured tunnels (see above). Higher layer security protection such as TCP-AO [RFC5926] is therefore optional, since it would be redundant with the security provided at lower layers.

Data communications over route optimized paths should employ end-to-end higher-layer security since only the control plane and unoptimized paths are protected by lower-layer security. End-to-end higher-layer security mechanisms include QUIC-TLS [RFC9001], TLS [RFC8446], DTLS [RFC6347], SSH [RFC4251], etc. applied in a manner outside the scope of this document.

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

10.1. Public Key Infrastructure (PKI) Considerations

In development of the overall ATN/IPS operational concept, ICAO addressed the security concerns in multiple ways to ensure coordination and consistency across the various groups. This also avoided potential duplicative work. Technical provisions related specifically to the operation of ATN/IPS are specified in supporting ATN/IPS standards. However, other considerations such as the establishment of a PKI, were determined to have an impact beyond ATN/IPS. ICAO created a Trust Framework Study Group (TFSG) to define various governance, policy, procedures and overall technical performance requirements for system connectivity and interoperability.

As part of their charter, the TSFG is specifically developing a concept of operations for a common aviation digital trust framework and principles to facilitate an interoperable secure, cyber resilient and seamless exchange of information in a digitally connected environment. They are also developing governance principles, policy, procedures and requirements for establishing digital identity for a global trust framework that will consider any exchange of information among users of the aviation ecosystem, and to promote these concepts with all relevant stakeholders.

ATN/IPS will take advantage of the developments of TFSG within the overall ATN/IPS operational concept. As such, this will include the usage of the PKI specification resulting from the TFSG.

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, Mach Chen, Russ Housley, Erik Kline, Hubert Kuenig, Tony Li, Gyan Mishra, Alexandre Petrescu, Dave Thaler, Pascal Thubert, Michael Tuxen, Eric Vyncke, Tony Whyman. Vaughn Maiolla is further remembered for his support and guidance.

Honoring life, liberty and the pursuit of happiness.

12. References

12.1. Normative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, DOI 10.17487/RFC0791, September 1981, <<https://www.rfc-editor.org/info/rfc791>>.
- [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.
- [I-D.ietf-lisp-rfc6830bis] Farinacci, D., Fuller, V., Meyer, D., Lewis, D., and A. Cabellos-Aparicio, "The Locator/ID Separation Protocol (LISP)", Work in Progress, Internet-Draft, draft-ietf-lisp-rfc6830bis-38, 7 May 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-lisp-rfc6830bis-38>>.
- [I-D.templin-6man-aero3] Templin, F. L., "Automatic Extended Route Optimization (AERO)", Work in Progress, Internet-Draft, draft-templin-6man-aero3-33, 3 March 2025, <<https://datatracker.ietf.org/api/v1/doc/document/draft-templin-6man-aero3/>>.

- [I-D.templin-6man-mla]
Templin, F., "IPv6 Addresses for Ad Hoc Networks", Work in Progress, Internet-Draft, draft-templin-6man-mla-25, 24 September 2024, <<https://datatracker.ietf.org/doc/html/draft-templin-6man-mla-25>>.
- [I-D.templin-6man-omni3]
Templin, F. L., "Transmission of IP Packets over Overlay Multilink Network (OMNI) Interfaces", Work in Progress, Internet-Draft, draft-templin-6man-omni3-40, 15 March 2025, <<https://datatracker.ietf.org/api/v1/doc/document/draft-templin-6man-omni3/>>.
- [IEN48] Cerf, V., "The Catenet Model For Internetworking, <https://www.rfc-editor.org/ien/ien48.txt>", July 1978.
- [IEN48-2] Cerf, V., "The Catenet Model For Internetworking (with figures), <http://www.postel.org/ien/pdf/ien048.pdf>", July 1978.
- [RFC4251] Ylonen, T. and C. Lonvick, Ed., "The Secure Shell (SSH) Protocol Architecture", RFC 4251, DOI 10.17487/RFC4251, January 2006, <<https://www.rfc-editor.org/info/rfc4251>>.
- [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>>.
- [RFC5926] Lebovitz, G. and E. Rescorla, "Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)", RFC 5926, DOI 10.17487/RFC5926, June 2010, <<https://www.rfc-editor.org/info/rfc5926>>.
- [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>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [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>>.

- [RFC6996] Mitchell, J., "Autonomous System (AS) Reservation for Private Use", BCP 6, RFC 6996, DOI 10.17487/RFC6996, July 2013, <<https://www.rfc-editor.org/info/rfc6996>>.
- [RFC8446] Rescorla, E., "The Transport Layer Security (TLS) Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446, August 2018, <<https://www.rfc-editor.org/info/rfc8446>>.
- [RFC9001] Thomson, M., Ed. and S. Turner, Ed., "Using TLS to Secure QUIC", RFC 9001, DOI 10.17487/RFC9001, May 2021, <<https://www.rfc-editor.org/info/rfc9001>>.

Appendix A. BGP Convergence Considerations

Experimental evidence has shown that BGP convergence time required after an MNP 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 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. AERO/OMNI Spanning Tree Properties

The AERO/OMNI services establish an adaptation layer for the OSI reference model appearing as "the layer below the network layer but above the data link layer". The adaptation layer presents a virtual bridging service from the network layer's perspective and a BGP-based IPv6 routing service from the data link layer's perspective.

AERO/OMNI overlay networks include s-ASBRs (aka Proxy/Servers) and c-ASBRs (aka Gateways) as the vertices in a graph with a typically sparse collection of edges between pairwise vertices. The graph minimally comprises a true spanning tree connecting all vertices but may also include additional edges in the spirit of IEEE 802.1aq Shortest Path Bridging. BGP routing then ensures loop freedom even if this "augmented" spanning tree includes cycles.

Each edge in the spanning tree corresponds to one or more connecting links which may be physical (e.g., fiber/free-space optics, etc.) or virtual (an IP tunnel over an underlying Internetwork). The OMNI interface provides a nexus for link selection, where control messages between vertices must be forwarded over edge links that are secured at the network layer or below while data messages may be forwarded over unsecured links. AERO/OMNI refer to these distinct forwarding contexts as the "secured spanning tree" and "unsecured spanning tree", respectively.

AERO route optimization provides an essential service to establish shortcut paths in the data plane that do not necessarily follow strict spanning tree paths. The shortcuts are formed through secured spanning tree control message exchanges which establish shortest path forwarding state in Clients, Proxy/Servers and intermediate Gateways. Route optimization therefore reduces spanning tree traffic concentration and instead distributes traffic over a diversity of underlying Internetwork paths.

Appendix C. Change Log

<< RFC Editor - remove prior to publication >>

Differences from -27 to -28:

- * Introduced MLAs and explained their role in the architecture.
- * Included reference to MLA spec.

Differences from -26 to -27:

- * Update references.

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

Gaurav Dawra
LinkedIn
United States of America
Email: gdawra.ietf@gmail.com

Acee Lindem
LabN Consulting LLC
United States of America
Email: acee.ietf@gmail.com

Victor Moreno
Cisco Systems, Inc.
United States of America
Email: vimoreno@cisco.com