

Network Working Group
Internet-Draft
Intended status: Informational
Expires: September 4, 2018

F. Templin, Ed.
G. Saccone
Boeing Research & Technology
G. Dawra
LinkedIn
A. Lindem
Cisco Systems, Inc.
March 03, 2018

A Simple BGP-based Mobile Routing System for the Aeronautical
Telecommunications Network
draft-templin-atn-bgp-06.txt

Abstract

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

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 September 4, 2018.

Internet-Draft

BGP for ATN/IPS

March 2018

Copyright Notice

Copyright (c) 2018 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Terminology	5
3.	ATN/IPS Routing System	6
4.	ATN/IPS Multilink and Mobility Service	9
5.	ATN/IPS Route Optimization	10
6.	BGP Protocol Considerations	13
7.	Implementation Status	14
8.	IANA Considerations	14
9.	Security Considerations	14
10.	Acknowledgements	15
11.	References	15
11.1.	Normative References	15
11.2.	Informative References	16
	Authors' Addresses	17

[1.](#) Introduction

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

The International Civil Aviation Organization [[ICAO](#)] is now undertaking the development of a next-generation replacement for ATN/

OSI known as Aeronautical Telecommunications Network with Internet Protocol Services (ATN/IPS). ATN/IPS will eventually provide an IPv6-based 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 a candidate 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 Traffic Engineering service that provides the highest possible degree of reliability.

The ATN/IPS is an IPv6-based [[RFC8200](#)] overlay network that assumes a worldwide connected Internetworking underlay for carrying tunneled ATM communications. The Internetworking underlay could be manifested as a private collection of long-haul backbone links (e.g., fiberoptics, copper, SATCOM, etc.) interconnected by high-performance networking gear such as bridges, switches, and routers. Such a private network would need to connect all ATN/IPS participants worldwide, and could therefore present a considerable cost for a large-scale deployment of new infrastructure. Alternatively, the ATN/IPS could be deployed as a secured overlay over the existing global public Internet. For example, ATN/IPS nodes could be deployed as part of an SD-WAN or an MPLS-WAN that rides over the public Internet via secured tunnels.

The ATN/IPS further assumes that each aircraft will receive an IPv6 Mobile Network Prefix (MNP) that accompanies the aircraft wherever it travels. ATCs and AOCs will likewise receive IPv6 prefixes, but they

would typically appear in static (not mobile) deployments such as air traffic control towers, airline headquarters, etc. Throughout the rest of this document, we therefore use the term "MNP" when discussing an IPv6 prefix that is delegated to any ATN/IPS end system, including ATCs, AOCs, and aircraft. We also use the term Mobility Service Prefix (MSP) to refer to an aggregated prefix assigned to the ATN/IPS by an Internet assigned numbers authority, and from which all MNPs are delegated (e.g., up to 2^{32} IPv6 /64 MNPs could be delegated from the MSP 2001:db8::/32).

Connexion By Boeing [CBB] was an early aviation mobile routing service based on dynamic updates in the global public Internet BGP

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

We therefore consider an approach using a BGP overlay network routing system where a private BGP routing protocol instance is maintained between ATN/IPS Autonomous System (AS) Border Routers (ASBRs). The private BGP instance does not interact with the native BGP routing system in the connected Internetworking underlay, and BGP updates are unidirectional from "stub" ASBRs (s-ASBRs) to a very small set of "core" ASBRs (c-ASBRs) in a hub-and-spokes topology. The Asymmetric Extended Route Optimization (AERO) architecture [[I-D.templin-aerolink](#)] is used to support mobility and route optimization services, where the BGP s-ASBRs are one and the same as AERO Servers and the BGP c-ASBRs are one and the same as AERO Relays. No extensions to the BGP protocol are necessary.

The s-ASBRs for each stub AS connect to a small number of c-ASBRs via dedicated high speed links and/or tunnels across the Internetworking underlay using industry-standard encapsulations (e.g., Generic Routing Encapsulation (GRE) [[RFC2784](#)], IPsec [[RFC4301](#)], etc.). The

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

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

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

[2.](#) Terminology

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

The terms "AERO Client", "AERO Proxy", "AERO Server", and "AERO Relay" are the same as defined in [[I-D.templin-aerolink](#)].

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

Air Traffic Managemnet (ATM)

The worldwide service for coordinating safe aviation operations.

Air Traffic Controller (ATC)

A government agent responsible for coordinating with aircraft within a defined operational region via voice and/or data Command and Control messaging.

Airline Operations Controller (AOC)

An airline agent responsible for tracking and coordinating with aircraft within their fleet.

Aeronautical Telecommunications Network with Internet Protocol Services (ATN/IPS)

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

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

Radio Access Network (RAN)

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

Core Autonomous System Border Router (c-ASBR) A BGP router located in the hub of a hub-and-spokes overlay network topology. Each c-ASBR is also an AERO Relay.

Stub Autonomous System Border Router (s-ASBR) A BGP router configured as a spoke in a hub-and-spokes overlay network topology. Each s-ASBR is also an AERO Server.

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

Proxy A node at the edge of a RAN that acts as a proxy go-between between Clients and Servers.

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 /64 MNPs could be delegated from the MSP 2001:db8::/32).

3. ATN/IPS Routing System

The proposed ATN/IPS routing system comprises a private BGP instance coordinated between ASBRs in an overlay network via tunnels over the Internetworking underlay (where the tunnels between neighboring ASBRs are set up as part of the BGP peering configuration.) The overlay does not interact with the native BGP routing system in the connected underlying Internetwork, and each c-ASBR advertises only a small and unchanging set of MSPs into the Internetworking underlay routing system instead of the full dynamically changing set of MNPs. (For example, when the Internetworking underlay is the global public Internet the c-ASBRs advertise the MSPs in the public BGP Internet routing system.) The routing system is discussed in detail in [\[I-D.templin-aerolink\]](#).

In a reference deployment, one or more s-ASBRs connect each stub AS to the overlay using a shared stub AS Number (ASN). Each s-ASBR further uses eBGP to peer with one or more c-ASBRs. All c-ASBRs are members of the same core AS, and use a shared core ASN. Since the private BGP instance is separate from the global public Internet BGP routing system, the ASBRs can use either a private ASN per [\[RFC6996\]](#) or simply use public ASNs noting that the ASNs may overlap with those already assigned in the Internet. (A third alternative would be to procure globally-unique public ASNs, but cost and maintenance requirements must be considered.)

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

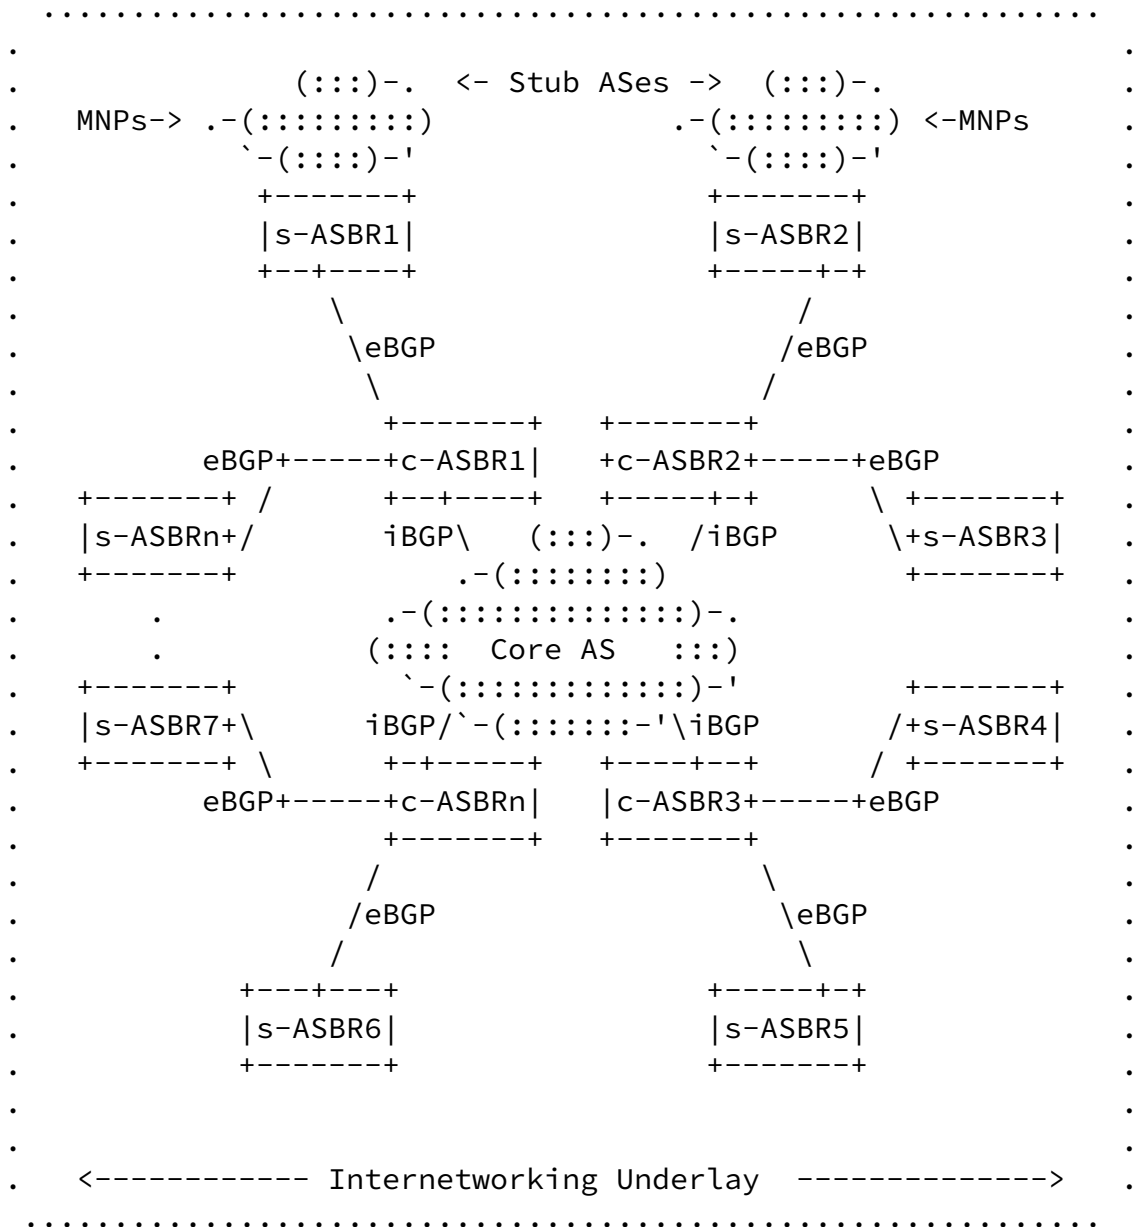


Figure 1: Reference Deployment

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

Since ATN/IPS end systems are expected to remain within the same stub

AS for extended timeframes, however, intra-domain mobility events (such as an aircraft handing off between cell towers) are handled within the stub AS instead of being propagated as inter-domain eBGP updates.

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

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

Therefore, assuming each c-ASBR can carry 1M or more routes, this means that at least 1M ATN/IPS end system MNPs can be serviced by a single set of c-ASBRs and that number could be further increased by using RRs. 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. For example, if the MSP for the ATN/IPS deployment is 2001:db8::/32, a first set of c-ASBRs could service the MSP segment 2001:db8::/40, a second set could service 2001:db8:0100::/40, a third set could service 2001:db8:0200::/40, etc.

In this way, each set of c-ASBRs services a specific set of MSPs that they inject into the Internetworking underlay native routing system, and each s-ASBR configures MSP-specific routes that list the correct set of c-ASBRs as next hops. This BGP routing design also allows for natural incremental deployment, and can support initial small-scale

deployments followed by dynamic deployment of additional ATN/IPS infrastructure elements without disturbing the already-deployed base.

4. ATN/IPS Multilink and Mobility Service

ATN/IPS end system multilink and mobility services are based on the AERO architecture [[I-D.templin-aerolink](#)], where end systems connect to aviation data link service provider Radio Access Networks (RANs). ATN/IPS end systems such as aircraft act as AERO Clients and may connect to multiple RANs at once, for example, when they have both a satellite link and an L-Band link activated simultaneously. Clients register all of their active data link connections with one or more AERO Servers which also act as s-ASBRs as discussed in [Section 3](#). Clients may connect to Servers either directly, or via an AERO Proxy at the edge of the RAN. The Proxy function corresponds to the manner in which web proxies communicate with web servers on behalf of clients in secured domains such as corporate enterprise networks.

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

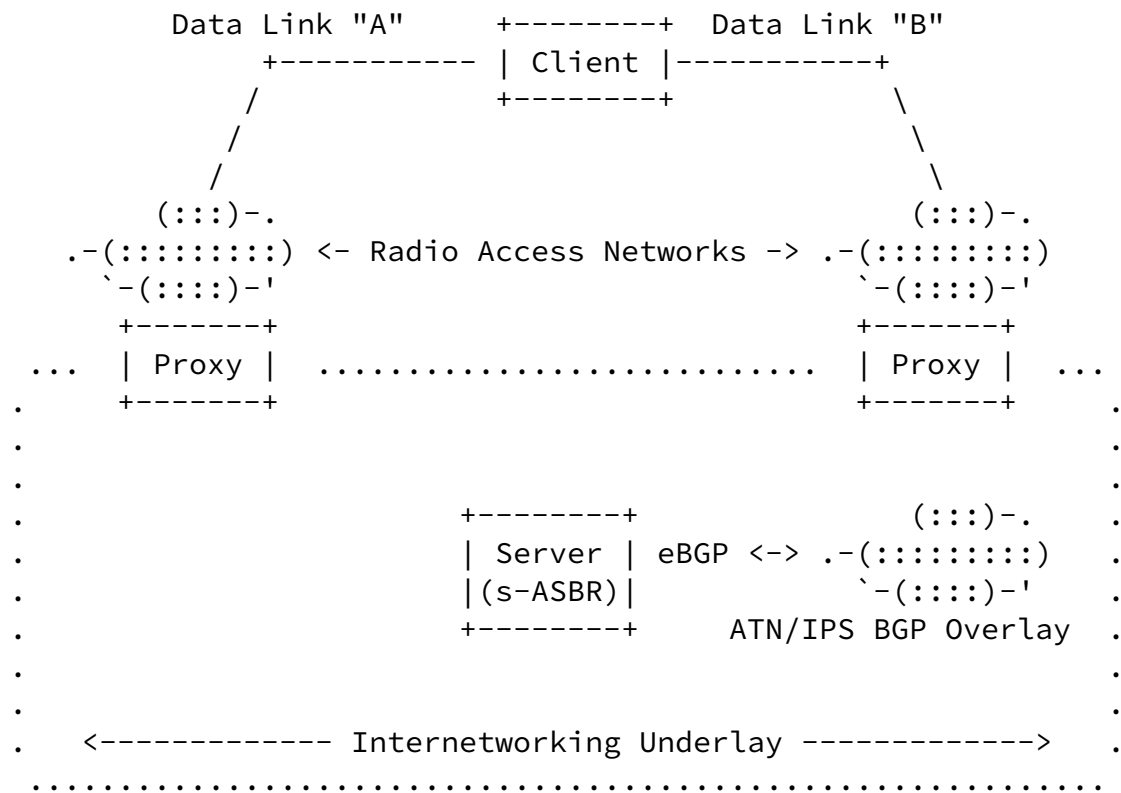


Figure 2: ATN/IPS Multilink and Mobility Architecture

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

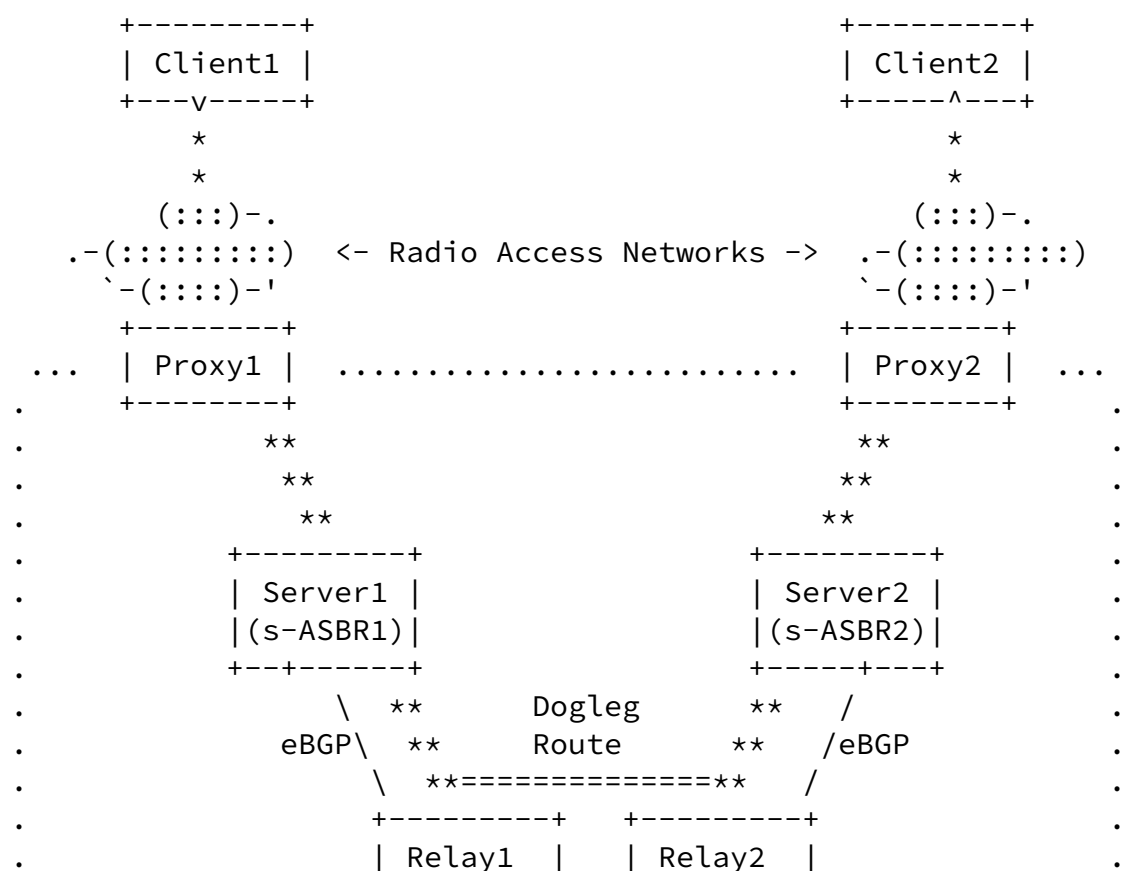
Once the Client has registered its data link addresses with the Server via one or more Proxies, the Proxies can signal fine-grained events like QoS changes to the Server on behalf of the Clients. For example, if a data link signal is fading, the Proxy can inform the Server without involvement of the Client. Moreover, if the RAN supports intra-domain route injection, the Client can avoid encapsulation and send and receive all of its packets unencapsulated since the RAN will natively route them to and from the Proxy. The Proxy will then tunnel the packets to and from the Server across the Internetworking underlay so that the Client need not incur any over-the-air encapsulation on performance-constrained aviation data links.

The Server represents all of its active Clients as MNP routes in the ATN/IPS BGP routing system. The Server's stub AS therefore consists of the set of all of its active Clients. The Server injects the MNPs of its active Clients and withdraws the MNPs of its departed Clients via BGP updates to c-ASBRs. Since Clients are expected to remain associated with their current Servers 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 Server handovers for aircraft operations (see: [Section 6](#)). It is important to observe that fine-grained events such as Client mobility and QoS signaling are coordinated only by Proxies and Servers, and do not involve other ASBRs in the routing system. In this way, localized events are not propagated into the global 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 ASBR peering topology discussed in [Section 3](#), this can initially only be accommodated by including multiple ASBRs-to-ASBR tunnel segments in the forwarding path. In many cases, it would be desirable to eliminate extraneous ASBR tunnel segments from this "dogleg" route so that packets can traverse a minimum number of tunneling hops across the Internetworking underlay using the AERO route optimization service [[I-D.templin-aerolink](#)].

A route optimization example is shown in Figure 3 and Figure 4 below. In the first figure, packets sent from Client1 to Client2 are transmitted across the source RAN to Proxy1 without encapsulation. Proxy1 then tunnels the packets to Server 1 (s-ASBR1), which tunnels them to Relay 1 (c-ASBR1), which tunnels them to Relay2 (c-ASBR2), which tunnels them to Server2 (s-ASBR2), which finally tunnels them to Proxy2. In the second figure, the optimized route tunnels packets directly from Proxy1 to Proxy2 without involving the ASBRs.



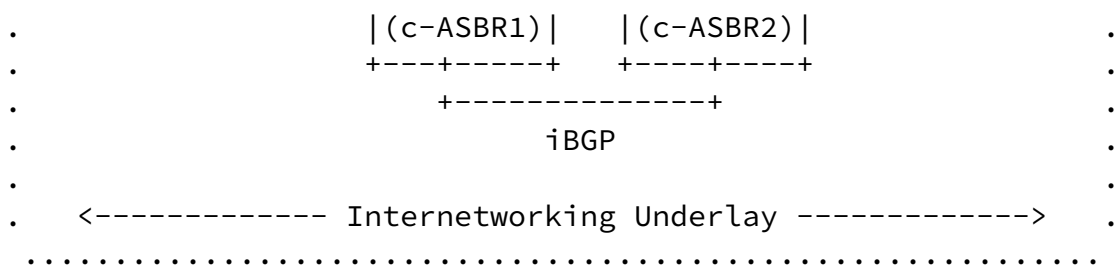



Figure 3: Dogleg Route Before Optimization






Figure 4: Optimized Route

The route optimization is accommodated by control message signaling between the Proxies and ASBRs. When the Proxy nearest the source sends a route optimization request, the request is forwarded toward the Server and nearest the destination. If the request is authentic, the destination Server provides the source Proxy with the address of the destination Proxy so that unnecessary tunnel segments are eliminated and direct Proxy-to-Proxy tunneling is enabled. At the same time, the destination Server keeps track of the source Proxies it has sent route optimization messages to so it can quickly update them if network mobility or Quality of Service (QoS) conditions change.

Note that route optimization can fail if Proxy1 cannot tunnel packets directly to Proxy2 due to some form of blockage in the Internetworking underlay such as filtering middle-boxes. It is also

necessary for Proxy1 to detect and adjust to failure of Proxy2 through receipt of a Server's IPv6 Neighbor Advertisement message and/or Neighbor Unreachability Detection (NUD) [[RFC4861](#)]. Note also that the Servers still maintain state so they can echo link QoS update messages coming from the RANs to inform correspondents of QoS changes (e.g., a link signal strength fading, a data link connection loss, etc.).

Finally, each s-ASBR always has a default route and can therefore always send packets via the dogleg route through a c-ASBR even if a route optimized path has been established. The direct paths between s-ASBRs and c-ASBRs are tunnels are maintained by BGP peering session keepalives such that, if a link or an ASBR goes down, BGP will detect the failure and readjust the routing tables. However, ASBRs and the links that interconnect them are expected to be secured as highly-available and fault tolerant critical infrastructure such that peering session failures should be extremely rare.

6. BGP Protocol Considerations

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

The number of aircraft in operation at a given time worldwide is likely to be significantly less than 1M, but we will assume this number for a worst-case analysis. Assuming a worst-case average 1 hour flight profile from gate-to-gate with 10 Server transitions per flight, the entire system will need to service at most 10M BGP updates per hour (2778 updates per second). This number is within the realm of the peak BGP update messaging seen in the global public Internet today [BGP2]. Assuming a BGP update message size of 100 bytes (800bits), the total amount of BGP control message traffic to a single c-ASBR will be less than 2.5Mbps which is a nominal rate for modern data links.

Industry standard BGP routers provide configurable parameters with conservative default values. For example, the default hold time is 90 seconds, the default keepalive time is 1/3 of the hold time, and the default MinRouteAdvertisementinterval is 30 seconds for eBGP

peers and 5 seconds for iBGP peers (see [Section 10 of \[RFC4271\]](#)). For the simple mobile routing system described herein, these parameters can and should be set to more aggressive values to support faster neighbor/link failure detection and faster routing protocol convergence times. For example, a hold time of 3 seconds and a MinRouteAdvertisementinterval of 0 seconds for both iBGP and eBGP.

C-ASBRs will be using EBGP both in the ATN/IPS and the Internetworking Underlay with the ATN/IPS unicast IPv6 routes resolving over Internetworking Underlay routes. Consequently, c-ASBRs and potentially s-ASBRs will need to support separate local

ASes for the two BGP routing domains and routing policy or assure routes are not propagated between the two BGP routing domains. From a conceptual and operational standpoint, the implementation should provide isolation between the two BGP routing domains (e.g., separate BGP instances).

[7.](#) Implementation Status

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

[8.](#) IANA Considerations

This document does not introduce any IANA considerations.

[9.](#) Security Considerations

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

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

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

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

[10.](#) Acknowledgements

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

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

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

11. References

11.1. Normative References

- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", [RFC 4271](#), DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4443] Conta, A., Deering, S., and M. Gupta, Ed., "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", STD 89, [RFC 4443](#), DOI 10.17487/RFC4443, March 2006, <<https://www.rfc-editor.org/info/rfc4443>>.
- [RFC4451] McPherson, D. and V. Gill, "BGP MULTI_EXIT_DISC (MED) Considerations", [RFC 4451](#), DOI 10.17487/RFC4451, March 2006, <<https://www.rfc-editor.org/info/rfc4451>>.
- [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>>.
- [RFC4861] Narten, T., Nordmark, E., Simpson, W., and H. Soliman, "Neighbor Discovery for IP version 6 (IPv6)", [RFC 4861](#), DOI 10.17487/RFC4861, September 2007, <<https://www.rfc-editor.org/info/rfc4861>>.
- [RFC8200] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", STD 86, [RFC 8200](#), DOI 10.17487/RFC8200, July 2017, <<https://www.rfc-editor.org/info/rfc8200>>.

11.2. Informative References

- [BGP] Huston, G., "BGP in 2015, <http://potaroo.net>", January 2016.
- [BGP2] Huston, G., "BGP Instability Report, <http://bgpupdates.potaroo.net/instability/bgpupd.html>", May 2017.
- [CBB] Dul, A., "Global IP Network Mobility using Border Gateway Protocol (BGP), http://www.quark.net/docs/Global_IP_Network_Mobility_using_BGP.pdf", March 2006.
- [I-D.ietf-lisp-ddt] Fuller, V., Lewis, D., Ermagan, V., Jain, A., and A. Smirnov, "LISP Delegated Database Tree", [draft-ietf-lisp-ddt-09](#) (work in progress), January 2017.
- [I-D.templin-aerolink] Templin, F., "Asymmetric Extended Route Optimization (AERO)", [draft-templin-aerolink-81](#) (work in progress), February 2018.
- [ICA0] ICAO, I., "<http://www.icao.int/Pages/default.aspx>", February 2017.
- [RFC2784] Farinacci, D., Li, T., Hanks, S., Meyer, D., and P. Traina, "Generic Routing Encapsulation (GRE)", [RFC 2784](#), DOI 10.17487/RFC2784, March 2000, <<https://www.rfc-editor.org/info/rfc2784>>.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", [RFC 4301](#), DOI 10.17487/RFC4301, December 2005, <<https://www.rfc-editor.org/info/rfc4301>>.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", [RFC 5246](#), DOI 10.17487/RFC5246, August 2008, <<https://www.rfc-editor.org/info/rfc5246>>.
- [RFC6836] Fuller, V., Farinacci, D., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol Alternative Logical Topology (LISP+ALT)", [RFC 6836](#), DOI 10.17487/RFC6836, January 2013, <<https://www.rfc-editor.org/info/rfc6836>>.
- [RFC6996] Mitchell, J., "Autonomous System (AS) Reservation for Private Use", [BCP 6](#), [RFC 6996](#), DOI 10.17487/RFC6996, July

Authors' Addresses

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

Email: fltemplin@acm.org

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

Email: gregory.t.saccone@boeing.com

Gaurav Dawra
LinkedIn
USA

Email: gdawra.ietf@gmail.com

Acee Lindem
Cisco Systems, Inc.
USA

Email: acee@cisco.com

Templin, et al.

Expires September 4, 2018

[Page 17]