Approved for Public Release; Distribution Unlimited Case # 07-0746

A ROUTING ARCHITECTURE FOR THE AIRBORNE NETWORK

Steven V. Pizzi

The MITRE Corporation
202 Burlington Road
Bedford, MA 01730

ABSTRACT

In this paper we present a candidate routing architecture for the future IP-based Airborne Network (AN). The advantages and disadvantages of this architecture are presented. We focus on the issues of mobility and the separation of routing domains.

The future AN will consist of multiple IP-enabled airborne platforms with wireless connections to each other and to multiple surface network domains. For example, Air Force platforms may connect to each other and may connect directly to Air Force domain ground sites. These Air Force platforms may wish to use Navy platforms as relays to access Navy ground sites and connect via SIPRNET to Air Force ground sites. Similarly, the best path may be to connect via an Army ground site to gain access to the Air Force Intranet via SIPRNET. Given the dynamically changing topology and the bandwidth-limited channel conditions corresponding to airborne networking versus terrestrial networking, it will be critical to develop effective link access protocols, routing protocols, and management strategies which can accommodate the unique characteristics of the Airborne Network.

Various IP-enabled radios (e.g., TTNT) would be used to establish these connections. Military satellite links, as well as various commercial satellite links, such as INMARSAT and Iridium, would also be available. The specific IP-radios that will be available for both the air and the ground nodes of the AN certainly will influence the eventual choice of routing architecture. We consider some of these IP-radios as part of our candidate architecture to ensure that our analysis is consistent with the planned infrastructure.

I INTRODUCTION

In this paper we present a candidate routing architecture for the future Airborne Network (AN), which will be based on the Internet Protocol (IP). The advantages and disadvantages are presented. We identify the protocols required and indicate their limitations. The work here builds on the Joint Airborne Network Services Suite (JANSS), discussed in [1].

We note that in developing a candidate AN architecture, the specific IP-radios that will be available for both the air and the ground nodes of the AN certainly will influence the eventual choice of routing architecture. We consider some of these IP-radios as part of our candidate architecture to ensure that our analysis is consistent with the planned infrastructure.

II BACKGROUND

The Airborne Network (AN) will consist of IP-enabled network nodes implemented aboard airborne platforms that intercommunicate as part of the Department of Defense (DoD) Global Information Grid (GIG). The GIG itself consists of three basic segments: Space, Airborne, and Terrestrial. Figure 1 provides a simplistic view of the GIG.

Individual DoD aircraft are working towards this IP-based airborne networking capability primarily to enable connectivity to the Secret Internet Protocol Router Network (SIPRNET). An additional benefit of this IP networking capability will be a more flexible, evolvable, and interoperable means of exchanging information between platforms and systems via this common method of transport.

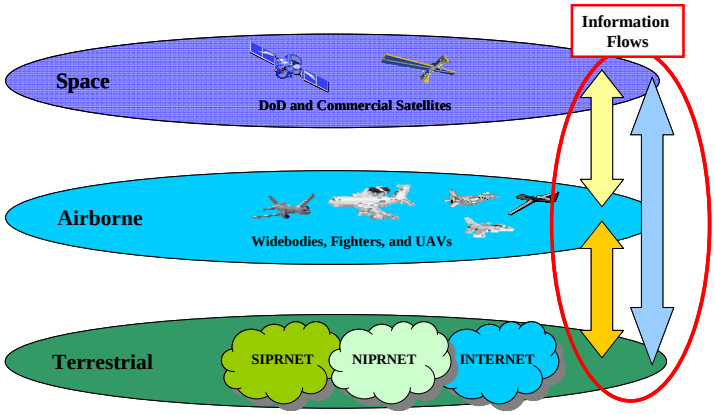

Figure 1. The Global Information Grid (GIG).

Unlike the infrastructure of the terrestrial network, the AN will be more dynamic. That is, an AN may exist for a given mission or series of missions in a specific

© 2007 The MITRE Corporation. All Rights Reserved.

geographical area. It would not necessarily be a permanent network over a particular area, since after a mission is completed the aircraft involved would land and the network would be dissolved. There may be different ANs corresponding to different theaters of operation, much like Joint Tactical Information Distribution System (JTIDS) networks.

Given the dynamically changing topology and the bandwidth-limited channel conditions corresponding to airborne networking versus terrestrial networking, it will be critical to develop effective link access protocols, routing protocols, and management strategies which can accommodate the unique characteristics of the Airborne Network.

Figure 2 shows the basic "vision" for the Airborne Network. Essentially, various aircraft operating within a theater would connect to each other via RF links and form an IP-based mobile routing network. Some of these aircraft would also connect to ground sites via satellite or radio links. Some of the satellite links may be commercial IP-networking services, such as INMARSAT or IRIDIUM. For DoD satellite IP-networking services the Transformational Satellite System (TSAT) program will provide a routed network among its satellite constellation [2].

As a starting point for AN routing architectures, we use the "vision" of Figure 2 as our basic topology and focus on implementing an IP-based AN using the currently planned terminals and systems.

Systems and Terminals

For this paper we consider the following emerging IP-based terminals, systems, platforms, and ground sites as elements in developing our candidate AN routing architecture:

1. Tactical Targeting Network Technologies (TTNT) System [3], [4];
2. Wideband Networking Waveform (WNW) System [4], [5];
3. Network Common Data Link (N-CDL) [6], [7];
4. INMARSAT Swift Broadband Service [8];
5. Joint Surveillance and Target Attack Radar System Aircraft (JSTARS, E-8);
6. Airborne Warning and Control System Aircraft (AWACS, E-3);
7. Combined Air Operations Center (CAOC);
8. Theater Network (A "Generic" Network within the local Theater of Operation);
9. Secret Internet Protocol Router Network (SIPRNET).

III ROUTING ARCHITECTURE

The key issues that we focus on in setting up an AN routing architecture are mobility and separation of administrative (routing) domains.

Figure 3 shows a simplified architecture for a widebody platform. As shown in this figure, a number of local area networks (LANs) could be supported on the airborne platform. These LANs access the network through the platform's AN Router, which may have multiple RF paths via IP-capable radios. The IP-radios would have their own "Mobile Network Routers," as indicated in the figure, which implement their respective mobile routing protocols, as part of their system. Since a widebody platform potentially would have the capability of installing more than one of these mobile network systems, it could serve as a gateway for these systems.

We focus herein on one basic routing architecture for the AN. Specifically, the architecture implements each of the mobile routing systems (i.e., TTNT, WNW, N-CDL) as a separate autonomous system (AS).

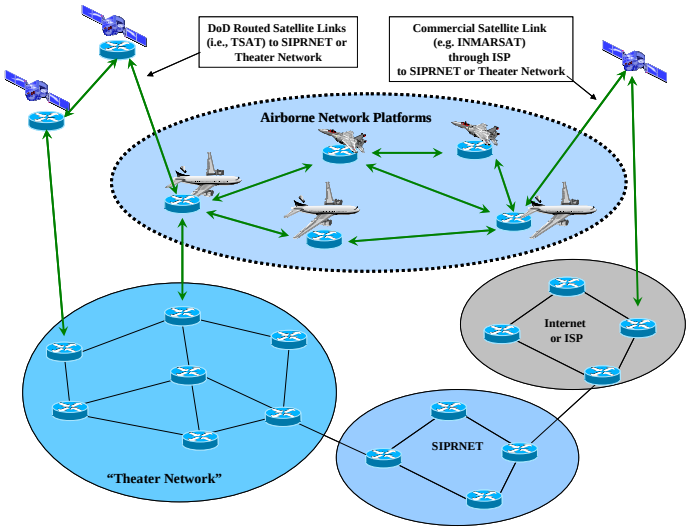

Figure 2. Airborne Network Vision.

An AS, or in alternative terminology an AS domain, is a separate administrative domain with its own routing policies [9]. An Internet Service Provider (ISP) would be one example of an AS domain. The network on a military base and the network of a small commercial company would be other examples of AS domains.

Candidate Architecture

Figure 4 represents an AN routing architecture where the mobile networks are separate individual AS domains.

© 2007 The MITRE Corporation. All Rights Reserved.

Here the JSTARS and the AWACS are shown as the widebody aircraft in the architecture. These connect to the SIPRNET and/or a "Theater Network" via INMARSAT. The Combined Air Operations Center (CAOC) is within the Theater Network. TTNT, WNW, and N-CDL are the mobile networks within this architecture. AWACS and JSTARS each have terminals (i.e., mobile router and IP radio combinations) that are part of these mobile networks. The CAOC also has a terminal that is part of the TTNT network.

In this architecture, the JSTARS, the AWACS, the CAOC, and the Theater Network are each separate autonomous system (AS) domains. JSTARS and AWACS can access the SIPRNET or the Theater Network via Virtual Private Network (VPN) IP tunnels [14] from the INMARSAT Ground Site, which is also shown here as a separate AS. Each of the mobile networks is also a separate AS.

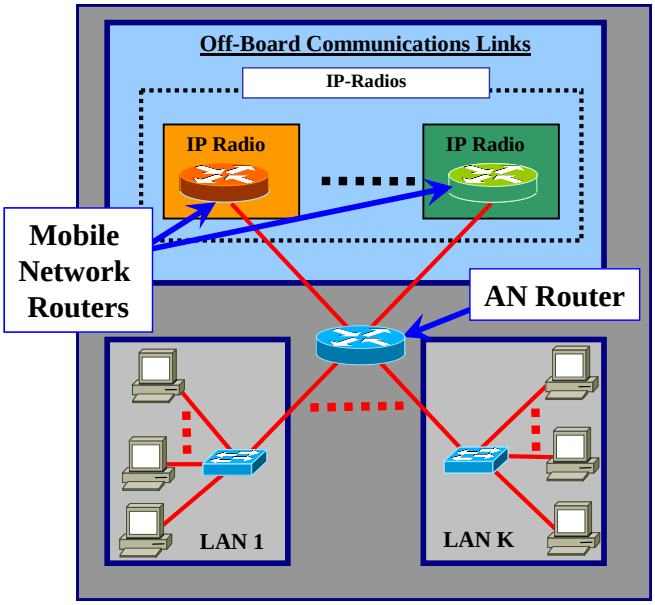

Figure 3. Simplified AN Platform Architecture.

For exchange of routing information between separate autonomous system (AS) domains (i.e., inter-domain routing), the Border Gateway Protocol is used [9], [10]. External BGP (EBGP) runs as a Transmission Control Protocol (TCP) session between the routers connecting separate AS domains. These routers are known as BGP border routers or gateway routers. Internal BGP (IBGP) runs as a TCP session between each BGP router pair within an AS. IBGP distributes within the associated AS domain the routes to other AS domains that are learned via EBGP.

Since JSTARS and AWACS are each separate AS domains, they need to run EBGP between their AN Routers and the Mobile Network Routers of the N-CDL, WNW, and TTNT AS domains. For INMARSAT, we show only an INMARSAT Radio; therefore we use the AN router for networking connectivity to the INMARSAT Ground Site.

As mentioned previously, IBGP would normally run as TCP sessions between BGP router pairs within an AS. As separate AS domains, the widebody platforms have only one gateway router node or border router (i.e., the AN Router). Therefore, IBGP would not need to run for the widebody platforms. Typically, IBGP would be implemented in terrestrial networks, which are more extensive and may have multiple AS-to-AS gateways.

The CAOC here is also run as a separate AS domain with EBGP running between its gateway router for the TTNT path and the path to the Theater Network.

Hosts/Routers within the Theater Network domain use routes generated by the Open Shortest Path First (OSPF) protocol [11], [12], [13] to connect to other hosts/routers. OSPF is also used for routing within the CAOC. Hosts/Routers within any of the AS domains would use BGP-generated routes to access the other AS domains.

In general, given the limited bandwidth of the various mobile networks, it would not be appropriate to use these networks for transit routing. That is, it probably would not be an efficient use of bandwidth to have an AS that is not a member of a mobile network be able to route through that mobile network to access another AS for general exchanges of information.

For example, with an INMARSAT link available to both JSTARS and AWACS, a host on the SIPRNET or within the Theater Network needing to access JSTARS would be able to do so directly via the VPN tunnels as shown by the dashed lines in Figure 4. We do not recommend that SIPRNET and Theater Network hosts be able to access JSTARS via the CAOC using TTNT. Therefore, EBGP from the CAOC AS would not advertise that route. That is, only the CAOC, JSTARS, and AWACS would be seen as members of the TTNT network and as gateways for other TTNT nodes.

The CAOC, as an element of the TTNT network, could communicate with JSTARS and AWACS via TTNT or INMARSAT. But given the limited bandwidth of this network, only "TTNT-specific" or "mission-specific" information should be sent between TTNT network members. Some form of differentiated services code point (DSCP) routing [15] may be useful here, to ensure that only higher priority traffic or "TTNT-specific" traffic is sent from the CAOC to JSTARS or AWACS via TTNT, while lower priority traffic is sent via the INMARSAT link.

© 2007 The MITRE Corporation. All Rights Reserved.

Information flows between JSTARS and AWACS may be similarly routed so that, for example, high priority traffic would flow over the TTNT path, medium priority traffic would flow via the WNW and N-CDL paths, and low priority traffic would be sent via the INMARSAT path.

Indeed there may be some mission-specific situations where a SIPRNET host would need to access JSTARS via the CAOC using TTNT. However, we expect these would be very specific situations and would not be the norm. As missions involving the IP-based AN evolve, these situations would need to be considered as part of mission/network planning.

Clearly, the appropriate concept of operations (CONOPS) for a given mission, the associated data that must be sent, and the available links for that mission will determine what options are available to efficiently send traffic between network nodes.

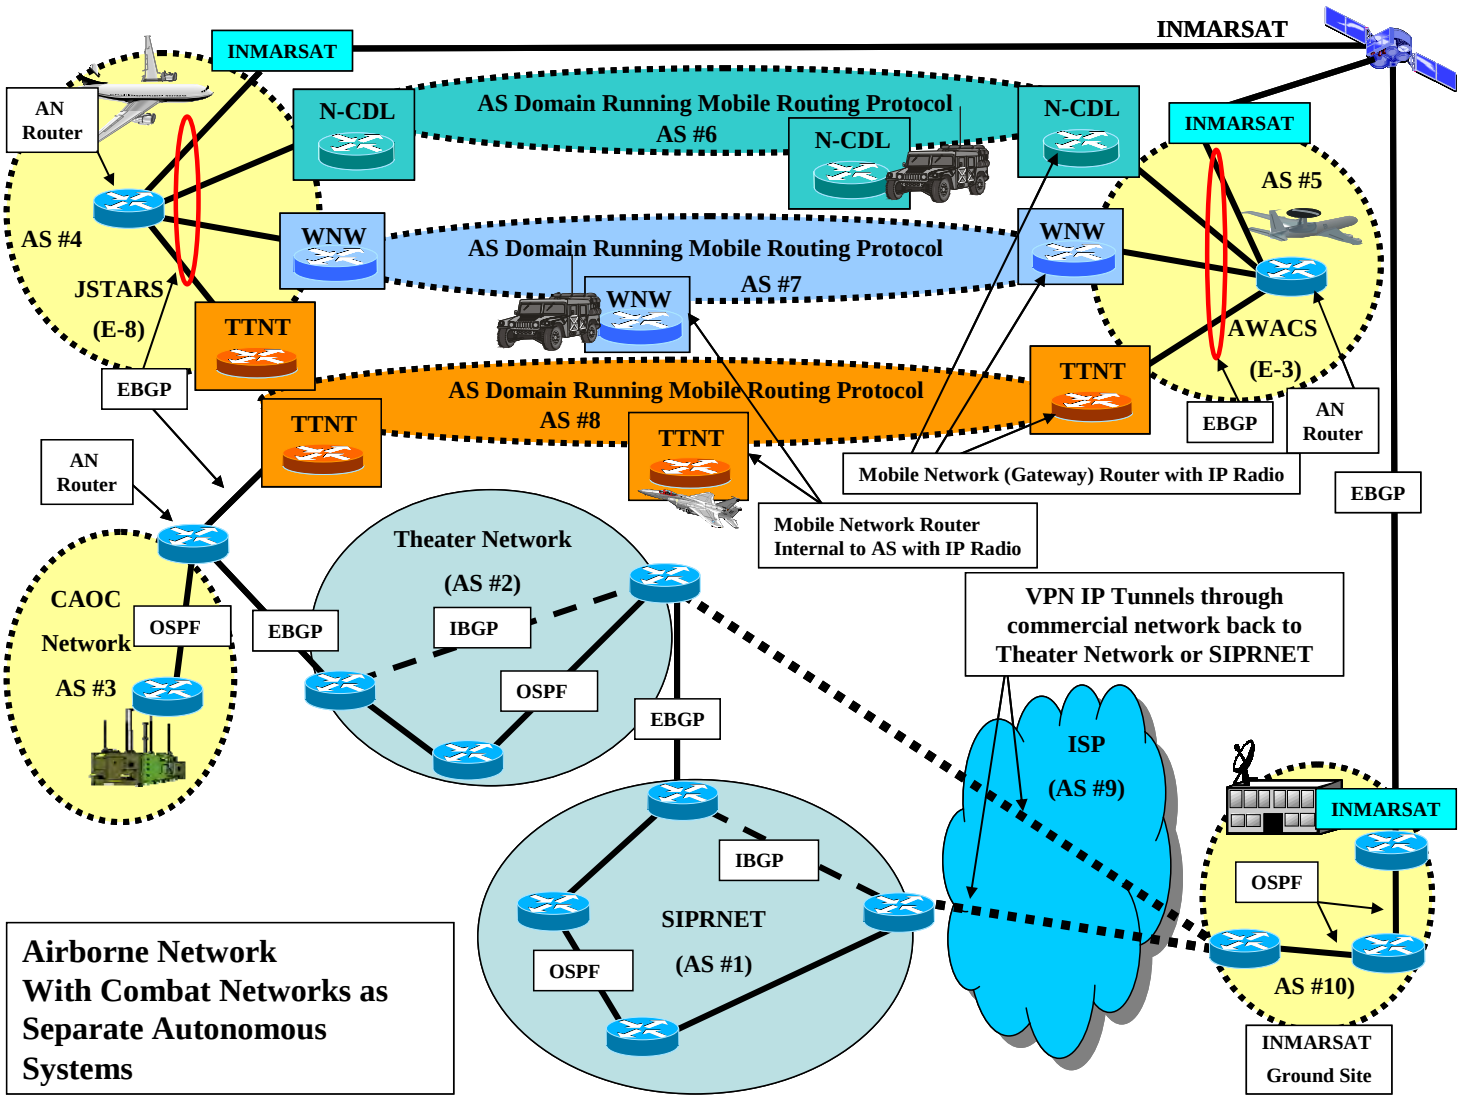

Figure 4. Airborne Network with Mobile Networks as Separate Autonomous Systems.

IV OTHER ARCHITECTURE OPTIONS

One of the key reasons for operating these different mobile networks as separate AS domains is to more easily facilitate adding-in nodes from other Joint Services. For

© 2007 The MITRE Corporation. All Rights Reserved.

example, if nodes from the Army's Warfighter Information Network - Tactical (WIN-T) system and the Navy's Automated Digital Network System (ADNS) wish to join, then BGP serves as the appropriate inter-domain routing protocol to connect these AS domains into the AN.

Nodes which are solely members of only one mobile network (e.g. fighter aircraft) would clearly operate only in that AS. These nodes may have a LAN on-board. However, they have only one path off the platform, which is via their Mobile Network Router/IP Radio.

As we have shown, some nodes which join these mobile networks may have membership in more than one of these AS domains. The CAOC, JSTARS, and AWACS of Figure 4 are the prime example of this. In cases such as these, it would be best to implement these nodes as separate AS domains via BGP. This is particularly important if these networks connect back to a larger infrastructure, such as the SIPRNET or a Theater Network, in order to prevent inadvertent transit routing between other AS domains over these mobile networks.

OSPF Options

An alternative to implementing these mobile networks as individual AS domains would be to implement these networks as some type of stub area within a larger OSPF network [16], [17]. For example, each mobile network could be implemented as an OSPF Totally Stub Area (TSA). The AWACS, JSTARS, and CAOC would be implemented as separate OSPF areas that are connected to these stub area mobile networks similar to Figure 4.

A TSA would receive no updates or be given any info regarding external destinations, since all external destinations would need to be reached via the TSA's Area Border Router (ABR) anyway [16], [17]. That is, only the default summary route (0.0.0.0 or "0/0") would be advertised by the ABR to the stub area [16]. The TSA option would limit any routing update overhead into the TSA. The reduction in overhead would be critical in limiting the amount of overhead on these wireless links, which would generally have less bandwidth than the terrestrial (typically optical fiber) links. The ABR serves effectively as a gateway for the TSA.

However, for OSPF we still need to define a backbone Area, defined as Area 0 [18]. All OSPF routers need to connect either physically or logically (i.e., via a tunnel) to Area 0. Suppose that the JSTARS platform network is defined as Area 0. Then suppose the AWACS platform network is defined as Area 1. Then since AWACS connects its Area 1 through the mobile network stub areas, the AWACS Area 1 requires a virtual path (i.e., a tunnel) to the JSTARS Area 0. Technically, in order to route to any of the other stub areas from its platform network Area 1, AWACS would have to send its packets to the backbone (Area 0) first via the Area 1-to-Area 0 tunnel. Then JSTARS would route those AWACS-originated packets to the appropriate TSA or to the INMARSAT link.

To avoid this tunneling to Area 0, static routes could be used. That is, specific routes (i.e., the IP addresses) would be manually configured into the AN routers of JSTARS and AWACS for the WNW, TTNT, and N-CDL networks. This is certainly acceptable, although it is less dynamic than OSPF. Another option would be to have the OSPF Area Border Routers discover shorter paths and avoid the virtual (tunnel) path [11].

However, we would not recommend this OSPF approach. The Figure 4 BGP approach is preferable. We note that to interface to INMARSAT's ground site, BGP may be required anyway. Therefore, since we would be configuring BGP for one link anyway, we may as well use BGP for all links.

RIP Options

Another option for the AN instead of OSPF and BGP would be to use the Routing Information Protocol (RIP) [19], [20], [21]. RIP does not rely on areas as does OSPF. RIP would be implemented between the AN Router and the individual Mobile Network Routers on each of the widebody platform networks. RIP would also be operated in a similar fashion for the CAOC. Since there are no "areas" in RIP, there is no need for the Area 0 tunnels, as in the previous OSPF option.

We would still need to run BGP between the CAOC and the Theater network to filter the TTNT routes in our example from the Theater Network to prevent transit routing. And we would still need BGP for the INMARSAT links to connect to that AS.

To implement RIP efficiently on the widebody platforms, route filtering would be done in RIP to prevent the entire AN Router routing table being advertised to each of the Mobile Network Routers. So, we could configure RIP to serve somewhat effectively as an alternative to BGP. But route filtering is essentially making the widebody platform serve as an AS anyway. Since BGP is the standard protocol for inter-domain routing; we recommend using BGP instead of forcing an intra-domain routing protocol perform the work of a true inter-domain routing protocol.

Further Discussion of BGP

BGP does not have a discovery process like the Hello Protocol of OSPF. TCP sessions between BGP router peers from the different AS domains need to be set-up manually. However as shown in Figure 4, the EBGP

© 2007 The MITRE Corporation. All Rights Reserved.

session are wired high-bandwidth connections between on-platform routers. Therefore, BGP configuration is under the individual platform's control with mobile routing information as supplied by the mobile network's management.

As mentioned previously, for typical BGP operation the three BGP routers within an AS would need to maintain an IBGP session between each pair. This was not necessary for the JSTARS or AWACS platforms, since there is only one BGP gateway router (i.e., the AN Router) for the AS on each of those platforms in the architecture of Figure 4.

The same IBGP pairing would be typically required for any of the Mobile Network Routers that are located within on the JSTARS and AWACS platforms or within the CAOC. For example, the TTNT Mobile Routers located on the JSTARS, AWACS, and within the CAOC run EBGP between themselves and the AN Router. Therefore, we might expect these three TTNT Mobile Routers to set-up IBGP sessions between each pairing. These IBGP sessions could be maintained over each individual mobile network. Or these sessions could be maintained via the INMARSAT link via virtual paths. Of course, given that these BGP routers are intended as gateways only, they would not really need IBGP to have the mobile AS domains learn about external routes. So most likely, the use of IBGP may be avoided for the mobile networks.

The TTNT, CAOC, AWACS, and JSTARS TTNT BGP routers, for example, would advertise their individual TTNT network IP address and their "gateway of last resort" address (0.0.0.0). TTNT terminals which need to send data out of the network would find the "closest" gateway. In this case there is probably no need for IBGP.

Clearly, these types of details need to be worked-out with the various mobile network developers.

Other Issues

So far we have assumed that the AN is a "red" network (i.e., operates at a single security level and that data protection is performed via link encryption). To further improve information assurance, we could add-in High Assurance Internet Protocol Encryptors (HAIPEs) to enable a "black core." "Black-core" issues involving the AN (e.g., [22], [23]) need to be addressed further and are beyond our scope. Other issues, such as implementing quality of service (QoS) and the use of performance enhancing proxies (PEPs) also need to be examined.

V CONCLUSIONS

In this paper we presented a candidate routing architecture for the Airborne Network. The primary characteristic of this architecture was to implement many of the networks and platforms as separate AS domains via BGP.

Specifically, we recommend the following be implemented as separate AS domains:

1. Mobile routing networks, such as TTNT, WNW, and N-CDL;

2. The widebody platforms, such as AWACS and JSTARS;

3. The CAOC and any other ground site which has multiple connections to other mobile networks and/or various terrestrial networks (e.g., Theater Network, SIPRNET).

This architecture allows for each widebody platform to be provided with its own set of unique IP addresses and be individually managed. The use of BGP allows for appropriate inter-domain address filtering and is the standard protocol used for inter-domain routing.

Each mobile routing network would be individually managed, much like JTIDS (Link-16) has its own management to configure the participating platforms.

By having the CAOC serve as a separate AS, transit routing via the mobile networks from the SIPRNET would be prevented via BGP.

An overall network planning organization, which would oversee and coordinate the allocation and distribution of IP addresses to each of the AN air and ground AS domains along with frequency planning and other Joint operational issues, is also needed.

Further details need to be addressed with platform integrators and terminal developers. The need for HAIPEs or other information assurance measures may dictate changes or modifications to this architecture.

VI ACKNOWLEDGEMENT

The author would like to thank D. Kenyon of the Air Force Electronic Systems Command, M. Solomon, K. Stranc, D. Dunbrack, R. Trafton, and M. Girard of the MITRE Corporation for their encouragement in this effort. The author would also like to thank D. Kiwior, D. Pecelli, E. Idhaw, B. Metcalf, J. Gauthier, S. McIntosh, and B. Savage also of the MITRE Corporation and G. Colella and M. Kraus of the Cisco Corporation for their helpful technical discussions.

VII REFERENCES

[1] R. Trafton and S.V. Pizzi, "The Joint Airborne Network Services Suite," MILCOM 2006, Washington., DC, October 2006.

© 2007 The MITRE Corporation. All Rights Reserved.
Case Number

[2] http://www.defenseindustrydaily.com/2005/07/special-report-the-usas-transformational-communications-satellite-system-tsat/index.php
[3] http://webext2.darpa.mil/body/news/2005/ttnt_05_demo.pdf
[4] http://integrator.hanscom.af.mil/2005/June/06162005/06162005-06.htm
[5] http://www.nova-eng.com/Inside.asp?n=Services&p=Wnw
[6] http://www.military-information-technology.com/article.cfm?DocID=1747
[7] http://www.military-information-technology.com/article.cfm?DocID=1976
[8] www.inmarsat.com
[9] B. Halabi and D. McPherson, Internet Routing Architectures, 2nd Edition, Cisco Press, New Riders Publishing, Indianapolis, IN, 2001.
[10] Y. Rekhter, "A Border Gateway Protocol (BGP-4), IETF RFC 1771, March 1995.
[11] J. Moy, "OSPF Version 2," IETF RFC 2328, April 1998.
[12] R. Coltun, D. Ferguson, and J. Moy, "OSPF for IPv6," IETF RFC 2740, December 1999.
[13] J. Moy, OSPF: Anatomy of a Routing Protocol, Addison Wesley Longman, Reading, MA, 1998.
[14] "Which VPN Solution Is Right For You," Document ID 14147, http://www.cisco.com/warp/public/707/which_vpn.html
[15] E. Crawley, R. Nair, B. Rajagopalan, and H. Sandick, "A Framework for QoS-Based Routing in the Internet," IETF RFC 2386, August 1998.
[16] "What Are OSPF Areas and Virtual Links," Document ID 13703, http://www.cisco.com/warp/public/104/8.html
[17] "How Does OSPF Generate Default Routes?" Document ID: 13692,http://www.cisco.com/warp/public/104/21.html
[18] "OSPF Design Guide," Document ID: 7039, http://www.cisco.com/warp/public/104/2.html
[19] C. Hedrick, "Routing Information Protocol," IETF RFC 1058, June 1988.
[20] G. Malkin, "RIP Version 2 Carrying Additional Information," IETF RFC 1723, November 1994.
[21] G. Malkin and R. Minnear "RIPng for IPv6," IETF RFC 2080, January 1997.
[22] G. Nakamoto, L. Higgins, and J. Richer, "Scalable HAIPE Discovery Using a DNS-Like Referral Model," MILCOM 2006, Washington, DC, October 2006.
[23] M. Mirhakkak, P. Ta, G. Comparetto, and V. Fineberg, "Modeling and Simulation of HAIPE," MILCOM 2006, Washington, DC, October 2006.
[24] http://www.isihellas.gr/mtps.html

© 2007 The MITRE Corporation. All Rights Reserved.