REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
OCTOBER 2009	Conference Paper Postprint	February 2009 – July 2009

4. TITLE AND SUBTITLE

RESERVATION-BASED QUALITY OF SERVICE (QOS) IN AN AIRBORNE NETWORK

5a. CONTRACT NUMBER
FA8750-09-C-0041

5b. GRANT NUMBER
N/A

5c. PROGRAM ELEMENT NUMBER
65502F

6. AUTHOR(S)

George Elmasry, Manoj Jain, Junghoon Lee, Roy Life, Gregory Hadynski, Bruce Metcalf

5d. PROJECT NUMBER
09SB

5e. TASK NUMBER
R0

5f. WORK UNIT NUMBER
15

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

XPRT Solution, Inc.
12 Christopher Way, Suite 301
Eatontown, NY 07724-1273

8. PERFORMING ORGANIZATION REPORT NUMBER
N/A

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

AFRL/RIGC
525 Brooks Road
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)
N/A

11. SPONSORING/MONITORING AGENCY REPORT NUMBER
AFRL-RI-RS-TP-2009-63

12. DISTRIBUTION AVAILABILITY STATEMENT
Approved for public release; distribution unlimited PA# 88ABW-2009-1461 Date Cleared: 10-August-2009

13. SUPPLEMENTARY NOTES
This paper was presented at and published in the Proc.of MILCOM 2009: Military Communications Conference 2009; Boston, MA, 18-21 October-2009. This work is copyrighted. One or more of the authors is a U.S. Government employee working, within the scope of their Government job; therefore, the U.S. Government is joint owner of the work and has the right to copy, distribute, and use the work. All other rights are reserved by the copyright owner.

14. ABSTRACT
This paper addresses the use of Resource ReSerVation Protocol-Aggregate (RSVP-AGG) at the tactical edge of the Air Force's Airborne Network (AN). Since the AN tactical edge can have different types of stub-networks accessing the AN (i.e., non-IP based legacy networks like Link 16, DiffServ based networks and IntServ based networks), RSVP-AGG offers a common access approach regardless of the differences in the networks using the AN. The paper presents a novel RSVP-AGG approach that has the advantage of decreasing the burden on the AN core links with limited bandwidth by reducing RSVP control traffic over the encrypted core. Also, the paper shows that RSVP-AGG (being a single reservation instead of multiple reservations) could be more resilient to link errors. Moreover,using RSVP-AGG over the AN core could open the door to consider the advantages of statistical multiplexing.

15. SUBJECT TERMS
airborne network, Integrated Services, IntServ, Quality of Service, QoS, RSVP, RSVP-AGG

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Gregory Hadynski
U	U	U	UU	8	19b. TELEPHONE NUMBER (*Include area code*) N/A

ID # 900129

RESERVATION-BASED QUALITY OF SERVICE (QOS) IN AN AIRBORNE NETWORK[1, 2]

George F. Elmasry, Manoj Jain, Junghoon Lee and Roy Life
DSCI, 12 Christopher Way, Eatontown, NJ 07724

Gregory Hadynski
AFRL/RIGC, 525 Brooks Road, Rome, NY 13441

Bruce Metcalf
The MITRE Corporation, 202 Burlington Rd., Bedford, MA 01730

ABSTRACT

This paper addresses the use of Resource ReSerVation Protocol-Aggregate (RSVP-AGG) at the tactical edge of the Air Force's Airborne Network (AN). Since the AN tactical edge can have different types of stub-networks accessing the AN (i.e., non-IP based legacy networks like Link 16, DiffServ based networks and IntServ based networks), RSVP-AGG offers a common access approach regardless of the differences in the networks using the AN. The paper presents a novel RSVP-AGG approach that has the advantage of decreasing the burden on the AN core links with limited bandwidth by reducing RSVP control traffic over the encrypted core. Also, the paper shows that RSVP-AGG (being a single reservation instead of multiple reservations) could be more resilient to link errors. Moreover, using RSVP-AGG over the AN core could open the door to consider the advantages of statistical multiplexing.

I. INTRODUCTION

The vision of the Air Force's Airborne Network (AN) is to interconnect tactical edge networks and offer reach-back capabilities for Warfighters. Future capability demands on the AN and its unique aspects create many challenges. This paper addresses the use of Resource ReSerVation Protocol-Aggregate (RSVP-AGG) between the enclaves of an AN. Since the edge enclaves can have different types and categories of traffic [i.e., non-Internet Protocol (IP) based legacy networks like Link 16, Differentiated Services (DiffServ) based networks such as Joint Tactical Radio System (JTRS), and Integrated Services (IntServ) based networks], RSVP-AGG offers a common access approach regardless of the differences in network enclaves using the AN. The paper also studies the effects of link reliability over RSVP and RSVP-AGG protocols, as well as the amount of savings of bandwidth (BW) for RSVP control signaling gained from using RSVP-AGG. Finally, RSVP-AGG over the AN core can lead to exploiting statistical multiplexing where multiple RSVP tunnels are replaced with a single tunnel. This is associated with an admission control mechanism that considers tactical networks' need for Multi-Level Precedence and Preemption (MLPP) where the higher precedence traffic is assured Quality of Service (QoS) at the expense of lower precedence traffic.

II. CHARACTERISTICS OF THE AN

Some of the assumptions for the AN that are different from those of commercial networks are as follows:

Bandwidth is Limited

Since the AN uses wireless links, spectrum availability can limit BW for the AN links. This is a major difference from wired networks, where fiber optic links typically offer a very large BW for applications requesting reservations, and most of the time, have a guaranteed, or even over-provisioned BW with virtually no error rate. Limited BW can mean that it is likely that a reservation will be denied. Hence, some intelligence is needed at the ingress and egress points to ensure that solving contention problems accommodates the mission's needs.

Traffic is Heterogeneous

While traffic in commercial networks is classified according to delivery requirements with classes of service such as interactive voice, interactive video, streaming voice, streaming video, data, ftp, e-mail, etc., traffic in tactical networks, such as the AN, can also be marked according to its precedence. Each class of service can have Routine, Priority, Immediate, Flash or Flash Override marking. Granting or denying a reservation should consider precedence requirements. For example, we should ensure that, we will not admit a Routine reservation, while we are denying a Priority reservation (for the same resources). We refer to this traffic prioritization as MLPP based IntServ. Another aspect of heterogeneous traffic is that some applications can be adaptable while some others are not. In other words, an application reserving a specific BW and getting a denial, can request a smaller BW, and if admitted, operate in a gracefully-degraded mode. It is anticipated

[1] Research conducted under SBIR AF073-020, "*Reservation-Based Quality of Service (QoS) in an Airborne Network*", sponsored by Air Force Research Labs, 525 Brooks Road, Rome, NY 13441.

that the AN would have a mix of applications, including those that can adapt and those that cannot.

Security is a Major Concern

The AN is expected to have a certain level of network security including encryption [1]. Thus, the RSVP approach [2] should consider secure reservation techniques where BW reservations are established through secure tunnels. With this assumption comes the need for the AN to have either pre-existing tunnels or to have a network manager that can plan and maintain the tunnels (based on the provision). We considered RFC 4804 [3], which allows the use of a dynamic, topology-aware admission control while reducing reservation signaling, which is in-line with the AN's needs. RFC 4804 also allows the use of a BW broker.

Subnets Accessing the AN use Different Technologies

The AN is a core network with different stub networks, or enclaves, communicating over it. Here are some examples of these stub networks:

- **Wired LAN in a C2 location**

  Access to the AN could be from a C2 (Command and Control) post with a standard wired Local Area Network (LAN). This LAN may have applications using standard RSVP, and the AN access point could offer RSVP-AGG to aggregate all these reservations under one tunnel [4].

- **Joint Tactical Radio System - Airborne, Maritime and Fixed (JTRS AMF) waveform**

  In this case, access to the AN is coming from a Differentiated Services Code Point (DSCP) based stub network enclave. The AN ingress point should be able to convert DiffServ to IntServ through techniques such as RSVP Proxy.

- **Legacy waveform (e.g., Link 16)**

  A legacy waveform may not be using IP at all, with a gateway converting its traffic into IP based packets. In this case, the ingress point should estimate the flow and create the necessary reservations over the AN.

Available Bandwidth may change due to Mobility

The RSVP-AGG or the RSVP Proxy technique at the AN ingress point needs to accommodate the changes in network capacity due to mobility of the terminals. Dynamic bandwidth changes in the AN backbone create very unique challenges for IntServ protocols. RSVP protocol needs to accommodate bandwidth fluctuations as well as an increased bit error rate that can result from mobility. Due to the constraints of encryption in the AN, monitoring capabilities such as Measurement Based Admission Control (MBAC) [5] are needed at the ingress and egress points (in addition to provisioning bandwidth) to detect the changes in the available bandwidth in the AN links in near real-time.

The Transmission Path may not be Bi-Directional

The RSVP-AGG or the RSVP Proxy technique needs to consider that we cannot always assume bi-directional transmission on the entire path. This could occur when a node is being jammed, causing it to transmit in one direction to a neighboring node and not receive from the same node along the same path. Thus, the node might need to form a uni-directional link to some other neighbor node to complete the circuit. This would create an issue for the use of RSVP-AGG in the AN since RSVP-AGG path messages, identical to RSVP, are generally bi-directional. Consequently, the application of any type of RSVP, including RSVP-AGG, offers a great challenge that is further addressed in the next section.

III. THE STUDY DETAILS

In this work, we focus on two aspects of the RSVP study for the AN. The first is to quantize the advantages of using RSVP-AGG over standard RSVP. The second is to quantize the effect of mobility on the RSVP protocol. (RSVP was originally designed for wired networks and its performance in wireless networks with packet loss can cause degradation in the protocol). Advantages of statistical multiplexing are currently under study.

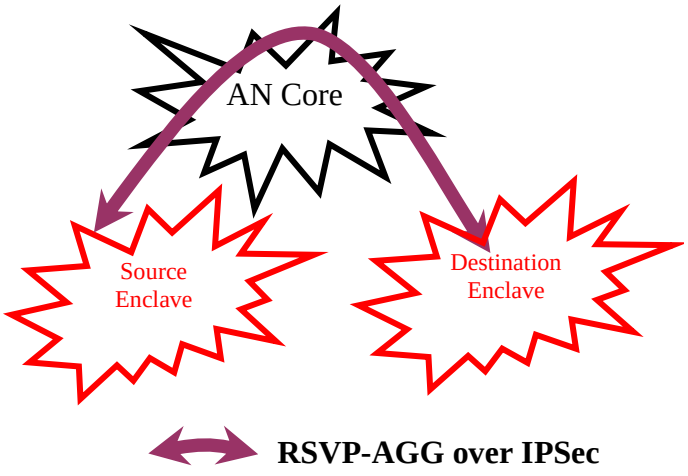

Figure 1: Red Enclaves Using Black AN Core for Connectivity

As Figure 1 shows, the AN core (where bandwidth utilization efficiency is needed), is envisioned to be encrypted [using Internet Protocol Security (IPSec) or Type 1 High

Assurance Internet Protocol Encryptor (HAIPE) encryption]. Red or plain text enclaves are envisioned to use the black AN core for internetworking connectivity. Secure RSVP-AGG tunnels are established between pairs of red enclaves over the black core. Conversion between RSVP flows in the red enclave and aggregated tunnels through the black AN core are performed by an edge device currently under development.

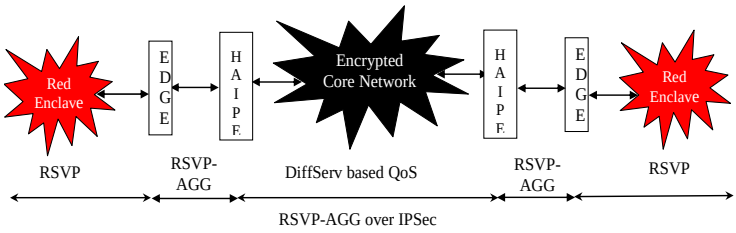

Figure 2: Implementation of RSVP-AGG over HAIPE

An implementation using HAIPE encryption can have problems. RSVP is based on end-to-end signaling and the current HAIPE specification does not allow RSVP signaling to be passed across encryption boundaries. As Figure 2 shows, individual RSVP flows originating in a red enclave will be aggregated by the (red) edge device and presented to the HAIPE with DSCP classification markings. The aggregated flows are then propagated through secure HAIPE tunnels, but only DiffServ priorities can be used through the black core. The benefit of guaranteed RSVP-AGG between enclaves over the black core is lost as a result of the HAIPE encryption. Even so, HAIPE v3.1 will allow the mapping of Explicit Congestion Notification (ECN) bits between the red and black sides of the encryption. These two ECN bits can be used as a component of a Measurement-Based Admission Control (MBAC) mechanism as they will provide the red edge device an indication of the level of black core network congestion [6]. Under conditions of congestion, an MLPP-based admission control process will be active in the edge device.

The current state of the art for handling RSVP between red enclaves over an encrypted tactical network core can be found in some of the WIN-T QED literature [5] where the QED proxies RSVP between red enclaves using MBAC. However, the QED approach does not use RSVP-AGG. This paper makes the case for considering RSVP-AGG instead of RSVP in order to: (1) Reduce control signaling over the encrypted core; (2) Be more tolerant to link error; and (3) Define a reasonable link reliability requirement to maintain RSVP signaling between the red enclaves.

The need for BW utilization efficiency that is driving this study leads to the consideration of creating a single RSVP-AGG tunnel (to accommodate all classes of service and all precedence levels) between each source and destination pair of red enclaves. As the results in the next section explain, creating a single tunnel can increase the AN core throughput efficiency since the control traffic needed to create and maintain multiple tunnels is significantly more than that for creating and maintaining a single, aggregated tunnel. Also, the single tunnel for RSVP-AGG is generated just once, whereas the standard RSVP tunnels are created, maintained for the life of each session, and then torn down at the end of the session.

Aggregation of RSVP at the AN ingress point and de-aggregation at the egress point will follow RFC 4804 [3], as mentioned before. Although creating a single tunnel for aggregated traffic increases the AN core throughput efficiency, it also generates a set of challenges regarding admission control. In other words, the aggregated traffic needs to be controlled at the AN ingress point using some MLPP policies as mentioned above. This, however, is beyond the scope of this paper and the reader can refer to [5] and [6] for more details about possible admission control policies in this encrypted environment.

In general, RSVP-AGG requires that the return path for RESV messages be identical to the forward path. However, if a bi-directional link is not available in the AN black core, the case needs to be looked at in terms of route discovery. The RSVP-AGG packets will be encrypted over the black core and will be routed according to the route tables in the black routers, which are based on Open Shortest Path First (OSPF). If a link is unidirectional due to jamming, etc., the routers will be discovered through this link in one direction, and the reverse path can be thorough different hops in a different link. In other words, on one hand, Type 1 HAIPE encryption offers a challenge in performing call admission control over the red enclave with limited information about the black core, while on the other hand, Type 1 HAIPE encryption offers us the ability to make RSVP-AGG (encrypted) messages pass over the black core even with uni-directional links.

IV. MODELING AN WITH RSVP-AGG

The study first started by using an OPNET model, which had some limitations since OPNET has not yet implemented the RSVP–AGG protocol. Thus, we relied on OPNET for multiple runs to study link reliability requirements, and built a testbed in order to study RSVP-AGG and quantify the amount of BW saving over RSVP in a fair manner. While the testbed enabled us to study RSVP-AGG accurately, OPNET model allowed us to run multiple scenarios under different link reliability (bit error rate) cases to create a comprehensive study.

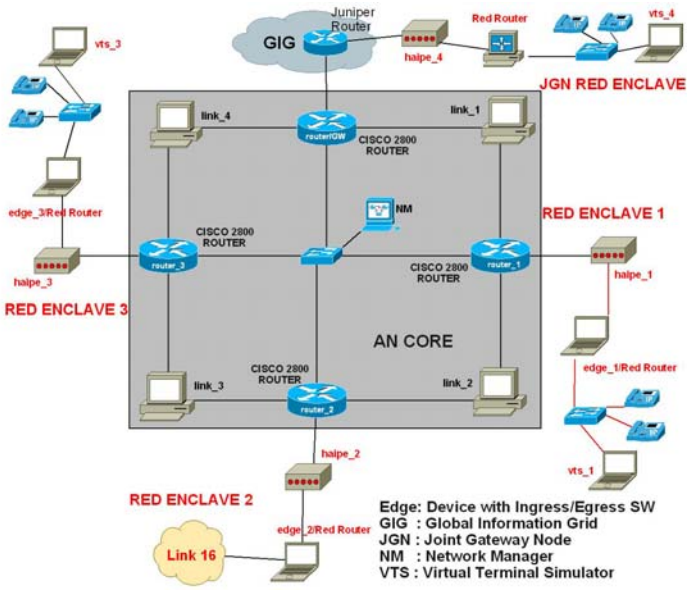

Figure 3: AN Testbed for RSVP-AGG Study

The AN testbed, shown in Figure 3, represents 4 different AN nodes with cipher text routers (Cisco 2800s). Also shown are four red enclaves (generating a mix of voice, video and data traffic) with red routers and HAIPE surrogates. A number of cases were run to study RSVP and RSVP-AGG. One case where two red enclaves had 10 simultaneous RSVP sessions was selected to compare RSVP-AGG control signaling against RSVP control signaling.

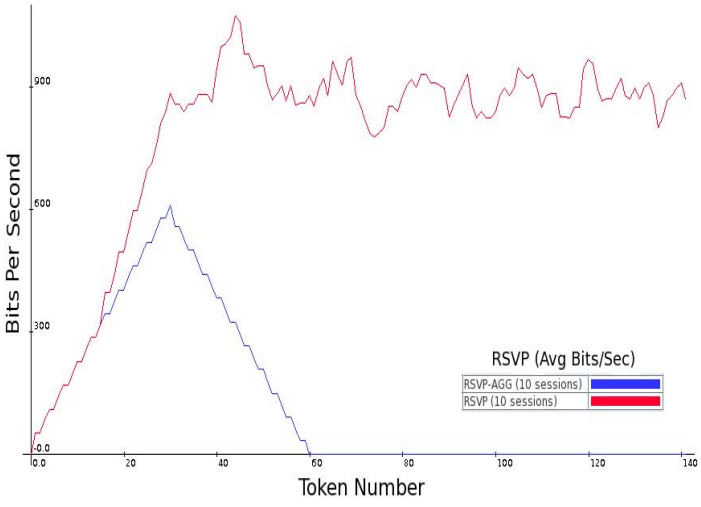

Figure 4: Control Signaling for RSVP-AGG vs. RSVP

Figure 4 shows the amount of measured control signaling in bps over the network core for the two cases. The horizontal axis refers to a token number which is a heart beat packet sent every 2 seconds between the two red enclaves and is used to piggyback the single RSVP-AGG maintenance packet. One can see how the RSVP-AGG maintenance traffic fades away while the RSVP maintenance traffic stays around 800-1000 bps. We had selected a configuration of RSVP that generates a maintenance packet every 30 seconds. More frequent maintenance packets will show a higher gain in BW savings.

V. BACKBONE NETWORK MOBILITY ANALYSIS

The OPNET model used to study link reliability represents a simple notional AN backbone model depicted in Figure 5. This model has a representation of the AN core through three aircrafts (JSTARS, AWACS, and Rivet-Joint). It should be noted the AN backbone representation of Figure 5 is purely notional and the only real significance to the aircrafts listed in the scenario is that those are all examples of widebody aircraft that could conceivably someday host AN backbone communications systems. There is no implication that we are simulating those platforms or that our results are in anyway indicative of the performance of their current communications systems.

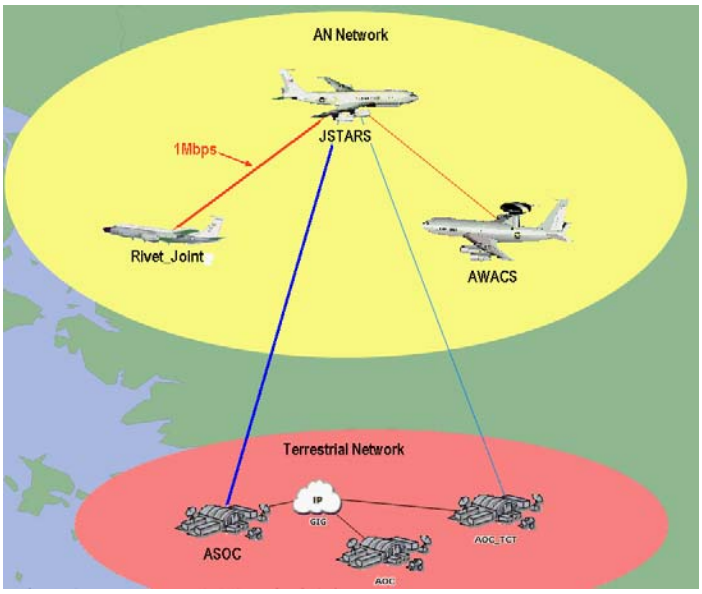

Figure 5: A Notional AN Backbone Model

In our model, a representation of a ground terrestrial network is used to inject external traffic into the AN core. In this study, we focused on the 1 Mbps link between the JSTARS and Rivet_Joint (shown in bold red in the figure) in order to create link reliability requirements.

Note that the RSVP-AGG scenario was created in OPNET as an approximation of the RSVP-AGG protocol which has not been implemented in OPNET. To simulate RSVP-

AGG, we created a single RSVP request at the edge that is equivalent to a bundle of sessions.

Mobility of AN backbone nodes can cause degradation in BER. (As the AN backbone nodes get further away from each other, the SNR of interconnecting radio links can start decreasing, resulting in degradation in the BER.) In this section, we analyze the relationship between the degradation in BER and the behavior of the RSVP mechanism.

In the OPNET model, an increase in the link BER causes the RSVP module to crash (the reservation terminates in an unexpected manner). As BER decreases, reservation becomes successful; however, it terminates (crashes) after a while (i.e., it cannot be kept or maintained). Initial results showed that the AN backbone links need to be maintained at a BER of 10^{-6} or better in order for the RSVP protocol to be established and maintained in OPNET. This raised the question: Since RSVP-AGG is anticipated to generate less traffic than the standard RSVP, can this translate into one being able to establish and maintain RSVP-AGG under a BER worse than 10^{-6}? To answer this, we relied on OPNET to run more what-if scenarios to compare the effect of mobility on RSVP-AGG versus its effect on the standard RSVP protocol.

Here we present link reliability results from three different scenarios:

- **Scenario 1** is the baseline scenario which studies the standard RSVP behavior when the mobility of the backbone causes degradation in the link (JSTARS to Rivet_Joint) being studied. A simulation of one hour was run for this scenario.
- **Scenario 2** replaces the standard RSVP with RSVP-AGG and produces comparable results. Two cases of this simulation were run, one for duration of one hour and the other lasting a day.
- **Scenario 3** adds additional traffic (from AWACS to Rivet_Joint, light red line in Figure 3) to produce results that could be compared to the other scenarios when the link under study (JSTARS to Rivet_Joint) shares both the initial traffic as well as the additional traffic. The three cases for the first two scenarios were repeated under this scenario.

The results of these simulations are presented below. Note that the results below are dependent on random number generation from OPNET. While using a common random variable approach [7] or a Monte Carlo processing [8] could yield results based on a larger sample, we found that the results presented below are sufficient to drive home the point that one needs a certain level of link reliability in order to maintain RSVP signaling between the red enclaves over the AN encrypted core.

Table 1: Scenario 1 – Baseline RSVP

Baseline Scenario - RSVP (1 Hr Simulation)	
BER	**Simulation Time for RSVP Crash**
0.1	0
0.01	0
1.0E-03	0
1.0E-04	577 sec
1.0E-05	1054 sec
1.0E-06	1hr – Completed, No Crash
0	1hr – Completed, No Crash

Table 1 shows the results obtained from the baseline scenario and lead us to the proposition that one needs to maintain a BER of 10^{-6} or better in order to make sure that RSVP works reliably. These results, however, used the standard RSVP which has many reservations going between the JSTARS and the Rivet_Joint nodes.

Table 2: Scenario 2 – RSVP-AGG for One Hour

RSVP-AGG (1 Hr Simulation)	
BER	**Simulation Time for RSVP Crash**
0.1	0
0.01	0
1.0E-03	1308 sec
1.0E-04	439 sec
1.0E-05	1hr – Completed, No Crash
1.0E-06	1hr – Completed, No Crash
0	1hr – Completed, No Crash

Table 2 shows the same set of results for the RSVP-AGG which has a single reservation. One can see that one hour of simulation was completed at BER of 10^{-5} (compared with the BER of 10^{-6} for the previous case). This can be explained simply based on the fact that there is less control traffic traversing the backbone network which would lead to fewer chances for a control packet to be in error that causes the protocol to crash. Note from the table that contrary to expectation, the simulation time for BER of 10^{-3} was longer than that for 10^{-4}. This could be just a coincidence, perhaps arising from a random error pattern that occurred earlier for the 10^{-4} BER case, affecting the protocol catastrophically.

Table 3: Scenario 2 – RSVP-AGG for One Day

RSVP-AGG (1 Day Simulation)	
BER	**Simulation Time for RSVP Crash**
0.1	0
0.01	0
1.0E-03	1757 sec
1.0E-04	720 sec
1.0E-05	1hr 57 min 56 sec
1.0E-06	23hrs 28min 28 sec
0	24 hrs – Completed, No Crash

The results in Table 2 prompted us to run the RSVP-AGG scenario for an entire day, with the results as presented in Table 3. This shows that at BER of 10^{-5}, the protocol crashed in a little less than 2 hours, while at BER of 10^{-6}, the protocol almost completed the entire day. This emphasizes the initial recommendation of keeping the BER at the AN backbone links at 10^{-6} or better.

All of the above results were conducted when the traffic was flowing between only two nodes (JSTARS and Rivet_ Joint), with the focus on the link between them. For Scenario 3, we added additional traffic flow from the AWACS to the Rivet_Joint node, in order to study the effect of mixed traffic on the performance of standard RSVP and RSVP-AGG protocols. The above three cases were simulated for the Scenario 3 and the results are presented in Tables 4, 5 and 6.

Table 4: Scenario 3 – Baseline RSVP

Baseline Scenario - RSVP (1 Hr Simulation)	
BER	**Simulation Time for RSVP Crash**
0.1	0
0.01	0
1.0E-03	604 sec
1.0E-04	695 sec
1.0E-05	2406 sec
1.0E-06	1hr – Completed, No Crash
0	1hr – Completed, No Crash

The results in Table 4 are similar to those in Table 1. From the table, one can see that for duration of 1 hour of simulation time, as before, one should maintain a BER of 10^{-6} or better in order to complete the scenario with the standard RSVP. The slightly longer simulation times (compared to Table 1) before the crash for the cases of BER of 10^{-3}, 10^{-4} and 10^{-5} are likely from random error patterns occurring somewhat later.

Table 5: Scenario 3 – RSVP-AGG for One Hour

RSVP-AGG (1 Hr Simulation)	
BER	**Simulation Time for RSVP Crash**
0.1	0
0.01	0
1.0E-03	625 sec
1.0E-04	586 sec
1.0E-05	1776 sec
1.0E-06	1hr – Completed, No Crash
0	1hr – Completed, No Crash

Table 5 repeats the same mixed traffic while using RSVP-AGG for one hour simulation between JSTARS and Rivet_Joint nodes. One can draw a conclusion similar to that from Table 4 that one needs to maintain a BER of 10^{-6} or better in order to complete the scenario simulation.

Table 6: Scenario 3 – RSVP-AGG for One Day

RSVP-AGG (1 Day Simulation)	
BER	**Simulation Time for RSVP Crash**
0.1	0
0.01	0
1.0E-03	625 sec
1.0E-04	586 sec
1.0E-05	1776 sec
1.0E-06	24hr – Completed, No Crash
0	24hr – Completed, No Crash

Table 6 shows the last set of results where we use RSVP-AGG with a mix of traffic and run the simulation for an entire day. One can see that similar to the conclusion from Table 5, one needs to maintain a BER of 10^{-6} or better in order to complete the scenario for the entire day.

VI. FUTURE WORK

This work is currently being extended to study the advantages of statistical multiplexing when all traffic is aggregated over a single tunnel instead of over multiple tunnels. As the tactical edge performs Admission Control (AC), it can admit more sessions in an aggregated tunnel than in separate tunnels (whereas the summation of the BW of the separate tunnels is equal to the aggregated tunnel BW and the QoS requirements in the two cases are the same). Note that the reservation request would be based not on the peak rate but the average rate. At a specific moment, some sessions in the aggregated tunnel would peak, i.e., use more bandwidth than the reserved, while other sessions might be using less than the reserved BW. Statistical multiplexing of the aggregated session can

create higher efficiency in using the AN backbone resources.

VII. SUMMARY

We ran a number of simulations of the performance of an airborne network using the OPNET model and a HW testbed, and showed the following:

1- RSVP-AGG has distinct advantages over the standard RSVP protocol in increasing the network's throughput. Reduced protocol overhead and the advantage of statistical multiplexing offered by ("longer-lasting") RSVP-AGG tunnels between AN enclaves result in significant bandwidth-utilization efficiency compared to individual ("short-term") RSVP tunnels per flow.

   RSVP-AGG bandwidth guarantees cannot be provided in a black core implementation with HAIPE enclave encryption. However, a measurement and MLPP-based admission control process can be implemented that exploits ECN bits mapped between the black core network and red enclaves (per HAIPE v3.1 specifications).

2- To maintain the RSVP-AGG protocol successfully, we must ensure that the BER over the AN backbone links does not exceed 10^{-6} (based on the OPNET simulation results, pending further investigations on the AN's behavior for transmission errors).

VIII. REFERENCES

[1] Junghoon Lee, George F. Elmasry and Manoj Jain, "Effect of Security Architecture on Cross-Layer Signaling in Network Centric Systems," *Proceedings of Milcom 2008*, NC9-3.

[2] Braden, R., Zhang, L., Berson, S., Herzog, S. and S. Jamin, "*Resource Reservation Protocol (RSVP) Version 1 Functional Specification*", RFC 2205, September 1997.

[3] F. Le Faucheur, "*Aggregation of Resource ReSerVation Protocol (RSVP) reservations over MPLS TE/DS-TE Tunnels*", RFC 4804, February 2007.

[4] F. Baker, et al, "*Aggregation of RSVP for IPv4 and IPv6 Reservations*", RFC 3175, September 2001.

[5] G. F. Elmasry, C.J. McCann, and R. Welsh, "*Partitioning QoS Management for Secure Tactical Wireless Ad-hoc Networks*," *IEEE Communications Magazine*, November 2005, pp. 116-123.

[6] George F. Elmasry, Junghoon Lee, Manoj Jain, Shane Snyder and Jonathan Santos, "*ECN-Based MBAC Algorithm for Use over HAIPE,"* forthcoming, MILCOM 2009.

[7] Sheldon M. Ross, "*Introduction to Probability and Statistics for Engineers and Scientists*", Fourth Edition, Academic Press, 2009.

[8] Alberto Leon-Garcia, "*Probability and Random Processes for Electrical Engineering"*, Second Edition, Addison-Wesley, 1994.