CHARACTERIZING TEST RANGE NETWORK INFRASTRUCTURE IN ANTICIPATION OF iNET DEPLOYMENT AND DESIGN

Joseph D Martin
Naval Air Systems Command
Patuxent River, MD USA

ABSTRACT

The iNET program uses network technology and infrastructure to enhance traditional telemetry systems. The program's components were designed with an eye to existing and emerging technology and infrastructure, requiring the program to gather data about these systems. The methods used in this design effort can be used to characterize existing network infrastructure to determine what upgrades and changes are necessary to deploy a TmNS. This paper describes the methods used for characterizing a range network infrastructure and explores network capacity and policy issues effecting a TmNS deployment. This effort includes making estimates and taking measurements of network capacity, surveying and analyzing network routing/management policies, and proposes a system for evaluating networks for future TmNS deployments.

KEY WORDS

iNET, TmNS, Infrastructure, Network Monitoring, Range Network

INTRODUCTION

In the process of developing a new standard for networked telemetry as a part of the integrated Network Enhanced Telemetry (iNET) program, it was necessary to gather a representative picture of existing range network infrastructure. This was done not only to determine a level of performance on which to base design decisions, but also to plot the first point on a roadmap for future improvements and upgrades to an existing range support network. This paper describes these efforts, including their goals, a test system architecture, test procedures, data collected, analysis of this data, and conclusions about the abilities of the range network tested with regard to networked telemetry.

INET BACKGROUND

The iNET program is working to complete a series of standards defining the Telemetry Network System (TmNS). TmNS uses existing network technologies to enhance traditional telemetry practices, with the goal of increasing flexibility and reliability in telemetry systems. By introducing these technologies, we gain the ability to communicate between any two or more elements of a telemetry system in a manner that is secure, reliable, and unimpaired by issues of vendor interoperability. Furthermore, Internet Protocol (IP) based network systems are already used by many ranges as a data transport mechanism between various geographically diverse locations. These existing networks will provide a basic ground infrastructure for TmNSs, lowering deployment cost, as well as taking advantage of existing network support resources

While incorporating network technologies into the world of telemetry has great potential, care must be taken to ensure that any proposed enhancements are within the abilities of the range network infrastructure on which they will be implemented. While many ranges have high speed and bandwidth IP networks already in existence, these networks have other duties and will only have a portion of their resources allocated to TmNS. Additionally, these networks may not provide consistent end-to-end performance. These concerns will be addressed below, with a particular focus on ensuring that an existing IP network architecture will support the design decisions made by the iNET program, and establishing a profile of the network resources that will likely be allocated to a TmNS.

DESCRIPTION OF ARON

The range network used to support this effort was the Atlantic Range Operational Network (ARON) located at Patuxent River Naval Air Station (Pax River NAS), in Patuxent River, Maryland. This network currently supports the Atlantic Test Range (ATR), transporting a variety of data, including telemetry, between ground station elements. Within the greater Patuxent River area, all network links are carried by fiber optic cable at a speed of 9621.5 Mbps (OC-192). However, some ATR antennas in locations which are not reachable by fiber, and are connected to the network via microwave link. These links are significantly slower than fiber links, and are also somewhat less reliable. This paper will pay specific attention to these links to ensure that they will support a TmNS.

The ARON is designed using a hub-and-spoke topology. The main hub is located in the main ATR range control facility. This facility includes both range control elements which manage the network, and Project Engineering Support (PES) rooms, which host range customers and allow them to view telemetry from their tests live. These facilities are implemented as spokes off the main hub. Additional spokes extend out to each antenna on the ATR. Some of these antennas are geographically close to the hub (across the street), while others are located much further away (some more than 50 miles). This paper focuses on one of these locations in particular: NASA's Wallops Flight Facility. This location was chosen because it represents a worst case scenario of low speed and bandwidth and less reliable microwave links. Two additional test locations were chosen for comparison, documented below.

GOALS

The main objective of this effort was to explore the ARON with regards to designing and implementing a TmNS. The data to be gathered can be grouped into three categories: Throughput, Reliability, and Policies. These three categories are explained below.

The throughput category includes various data points relating to the total amount of data that the ARON can transport between two nodes. Of interest are the total aggregate data rates that the ARON can sustain, as well as short term burst data rates. In addition, packet round trip time between two nodes was also measured.

The reliability category addresses if and for how long network links were unavailable. Fiber links were expected to be highly reliable, and only fail in the event of physical damage to the fiber itself. Microwave links however can be interfered with, and may suffer degraded performance due to external factors (atmospheric, rf, etc.). The failure of a network link to an antenna removes that antenna from the pool of antennas available to support a TmNS, and degrades the overall system effectiveness. Degraded system performance can also occur when a network is loaded at or close to capacity. To characterize this effect, the performance of various network links was studied with the ARON under both simulated and actual loads.

The policy category addresses whether any ARON network administrative or routing policies will impact TmNS deployment. A TmNS relies on many existing network technologies, including IP Multicast (Internet Group Messaging Protocol, or IGMP) and traffic prioritization/Quality of Service (DiffServ). While these technologies are widely supported, some networks are configured to disable these services as a security measure. To ensure the successful implementation of a TmNS, these technologies were exercised to ensure that they function as expected.

By gathering data in each of these three areas, we were able to create a representative network environment in which to continue to test and develop the standards that make up the TmNS. We also gathered data about the capabilities of an existing range network, and can now use that data to inform future design decisions about the TmNS.

TEST SYSTEM ARCHITECTURE

The test system consisted of several laptop computers deployed at various locations on the ARON. One primary laptop, acting as a Test and Monitoring Server (TMS), acted as a master coordinating all tests and recording data for analysis. Several additional laptops, acting as Remote Test Nodes (RTNs), acted as targets and partners for the TMS to use for running various tests. The TMS was deployed to the main ATR range control facility and network hub, so as to be at a central point on the network. The RTNs were deployed at various antenna sites around the ARON, primarily at NASA's Wallops Flight Facility (located approximately 55 miles from the main ATR range control facility) and Pax River NAS building 1591 (located approximately 2.5 miles from the main ATR range control facility). These laptops used a variety of software to monitor, measure, and record network performance metrics over the course of several weeks in July of 2011.

The computers used in this test were Dell Latitude E6500 laptops. Each was equipped with one Intel Core 2 Duo P8700 processor, running at 2.53 GHz. Each laptop had 4 GB of RAM, and one 250 GB hard drive. Each laptop was also equipped with a 1000-Base T Ethernet adapter. These laptops were connected to the ARON distribution switch at each location via Cat5e Ethernet cable. The switches used at each location were one of a Cisco 3560, 3750, or 2960.

Each laptop was configured to run a suite of Open Source and custom software developed for this test. Open Source Software was used where possible to facilitate test repetition at other ranges. The use of Open Source Software also allowed the use of standard network software tools such as NetSNMP and Nagios (see below).

Several Open Source Software projects were used in this test. Fedora Core 13 was used as an operating system. Fedora was chosen because it has a proven track record of stability; something necessary for deploying the RTNs where they would not be readily accessible (the antennas at Wallops Island are difficult to access). Fedora is based on on the widely popular Linux Kernel, and is compliant with all relevant unix standards (POSIX), allowing it to run all of the necessary software for this test. Additionally Fedora is freely redistributable, making it easy to use Fedora again if repeating this test at another range.

The Nagios Network Management System was used to coordinate the various tests run by the TMS and RTNs and to record the data from these tests. Nagios is widely used in industrial and academic environments, including use by AT&T, Google, Hewlet Packard and IBM. Nagios is primarily written in the perl programming language. Nagios' architecture is based on "plug-ins," which are also written in perl. This allows users to modify existing plug-ins or create new ones. One major advantage of Nagios is the Nagios Remote Plug-in Executor (NRPE), which provides a secure mechanism for the TMS to remotely execute tests from a RTN. This allows both sides of a link to be tested, to identify links with non-symmetrical characteristics. Another advantage of using Nagios is it's wide support community, which develops many useful add-on software projects. One such project is NagiosGraph, which creates highly flexible visualizations of test data automatically. NagiosGraph proved invaluable in processing and analyzing the data gathered in this test.

Several other Open Source Software projects were used in this test. NetSNMP is a standard library implementing the Simple Network Management Protocol, and was used by Nagios to manage various network devices. The File Transfer Protocol server "ftpd" and Hyper Text Transfer Protocol server "Apache" were both used to distribute files across the ARON, and to generate network traffic to simulate a load. The Firefox web browser was used as a front-end to access Nagios' web interface. The software package Emcast, developed by the University of Michigan, was used as a generic IP Multicast traffic generation tool. The network traffic monitoring tool Wireshark was used to monitor network traffic and verify the test system was functioning as desired.

A custom IP unicast traffic generation tool was created for this test. Several existing tools were evaluated, but none had a suitable feature set. The tool, "bwTester," was written in C++ and allows several tests to be scripted and then run without user intervention. It allows the user fine grain control over various traffic parameters, including the DiffServ priority level and the network traffic payload. Additionally, various scripts were written in perl and bash to automate some of the test functions.

TEST PROCEDURES

This test was conducted in three phases: a Test Development Phase, a Passive Monitoring phase, and an Active Testing Phase. These phases are described below.

In the Test Development phase, the ARON design was studied in detail to determine how best to characterize it's performance. Each link was categorized according to its link type and anticipated reliability and bandwidth. Test locations were then chosen for each category to create a full picture of the ARON. For each location, a set of metrics was identified to address each of the stated test goals.

For each metric, a procedure was established and software to measure the metric was identified. To measure throughput, a connection would be established between the TMS and a RTN. Data would be sent over this connection at various aggregate speeds, and the amount of data actually sent over a period of time was measured. Short term burst speeds were also measured in this way. To measure a link's round trip time, the "ping" utility was used. To measure link reliability, the "ping" utility was again used. Pings were sent every 30 seconds over the course of the test (several weeks), and unreturned pings were noted. To test various policy related metrics, custom IP packets were generated using the custom traffic generator, the "ping" utility, and Emcast. These packets were then monitored at another test machine with Wireshark to ensure they were passed through the network cleanly.

Next, the test laptops were set up and configured. All necessary software was installed on the TMS and one RTN. These machines were set up in a lab environment and were left running for a period of 48 hours to ensure they were stable. The RTN was then cloned to create as many individual RTNs as were needed. When all RTNs were configured, another 48 hour stability test was completed, this time with all RTNs present on the network.

The goal of the Passive Monitoring phase was to establish a baseline for various network performance metrics with minimal impact to the network. These tests were conducted throughout the day, including during periods when the range was active. The TMS was deployed to the main ATR range control facility and RTNs were deployed to Pax River NAS building 1591 and NASA's Wallops Flight Facility.

Every 30 seconds, Nagios would use the "ping" utility to send an IP packet from the TMS to each RTN and from each RTN to the TMS. When each ping returned to its origin, its round trip time as measured by the system's CPU clock was measured and recorded. If a ping was not returned within 10 seconds, it was assumed that the link was unavailable. All periods of link unavailability were noted. At the end of the Passive Monitoring phase, data was analyzed and a baseline for normal network performance was established.

The goal of the Active Testing phase was to measure network performance while placing a load on the network representative of one or more TmNS implementations. Based on early versions of the iNET standards, it was assumed that a TmNS implementation with one active test would require approximately 20 Mbps of sustained network bandwidth with significantly higher asynchronous bursts. Connections were established between the TMS and various RTNs at bandwidths ranging from 20 Mbps (emulating a single test within a single TmNS) up to 500 Mbps (emulating multiple TmNSs with multiple tests at burst speeds). The software that created these connections was configured to log any periods where the network was unable to maintain the desired bandwidth. While these connections were established, the tests from the Passive Monitoring phase were repeated and their results recorded for later analysis and comparison with the resting baseline.

The Active Testing phase also included the policy related tests. To test that the ARON would support DiffServ, custom IP packets were created, exercising the full range of DiffServ Code Point (DSCP) values (in DiffServ, each DSCP values corresponds to a priority level). These packets were created both by the “ping” utility and the custom traffic generation tool used for the above load testing. At the packet’s destination, a computer running Wireshark captured the packets and checked them to verify that the packet’s DSCP value was unaltered. The ARON’s ability to support multicast was tested in a similar way. Emcast was used to create IGMP connections between two or more of the test machines, and data was sent to and from each machine to ensure that these connections could be established and maintained without issue. These connections were also analyzed with Wireshark to ensure they were behaving as expected.

To avoid impacting the day to day operations of the ATR the Active Testing phase tests were all conducted at night during the ARON “maintenance window” from 2300 - 0500.

DATA

Link throughput was measured according to two metrics: sustainable link bandwidth and link round trip time. For all fiber optic links tested, the sustainable link bandwidth was found to be above 500 Mbps. Links could be opened at this speed with a negligible initial drop in link latency (< 4 ms). Once these links were established, overall latency through the link actually dropped, due to the link remaining established for a period of time. The average resting round trip time for these links was measured at 867.38 µs, with an isolated peak resting round trip time at 1.44 ms (Figure 1). When loaded the round trip time peaked at 1.73 ms, but the average round trip time stayed consistent at 851.42 µs (Figure 2).

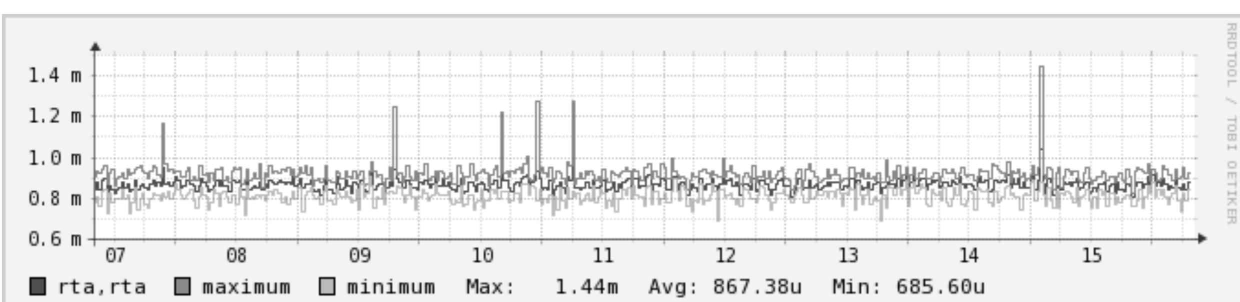

Figure 1 – Ping round trip time over the course of one week between the ATR Range Control Facility and a local antenna site over a fiber optic network link

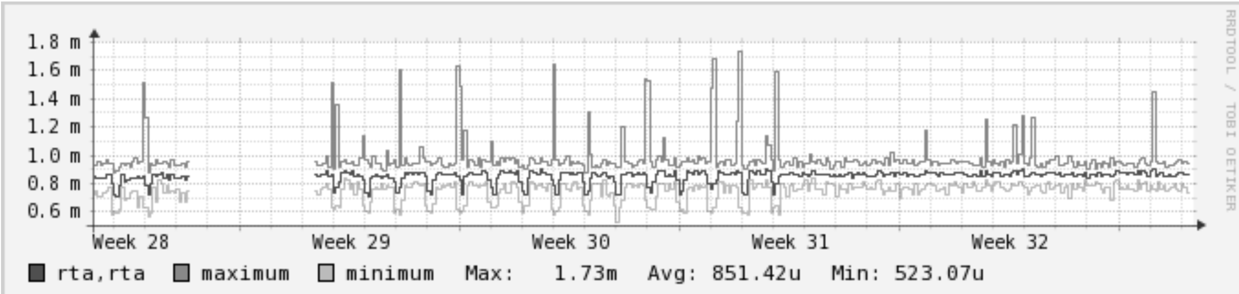

Figure 2 – Ping round trip time over the course of several weeks between the ATR Range Control Facility and a local antenna site over a fiber optic network link showing active testing

For the NASA Wallops Flight Facility microwave link, the sustainable link bandwidth was found to be roughly symmetrical and variable within the range of 36.50 to 43.61 Mbps from the Wallops Flight Facility and 36.14 to 42.20 Mbps to the Wallops Flight Facility (Figure 3). When operated at these loads, the average link round trip time was 150.53 ms with a peak at 300.34 ms (Figure 5). The average resting round trip time for this link was measured at 2.09 ms with a peak at 2.34 ms during the test period (Figure 4; see Reliability section).

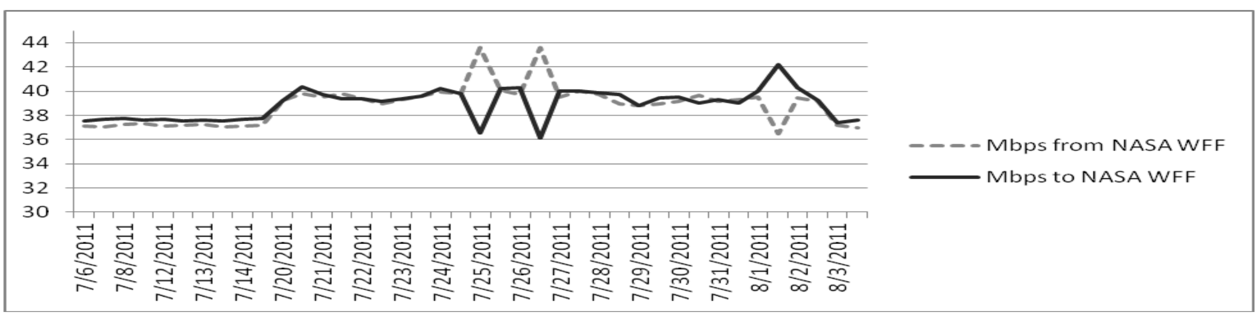

Figure 3 – Link throughput over the course of several weeks between the ATR Range Control Facility and the NASA Wallops Flight Facility over a microwave network link

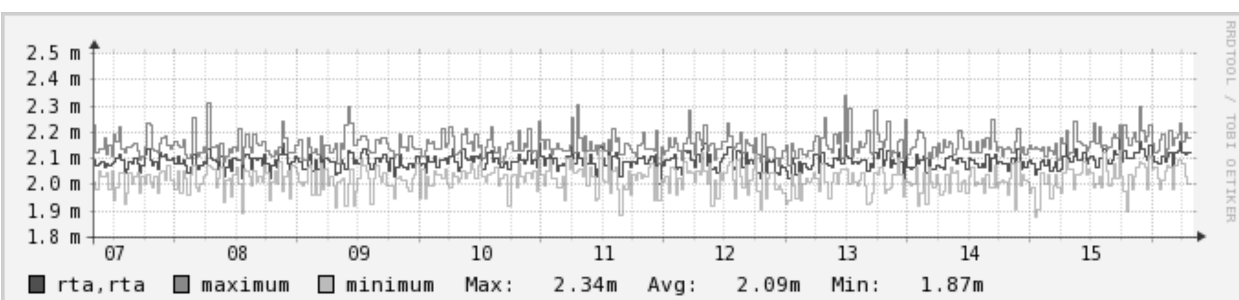

Figure 4 – Ping round trip time over the course of one week between the ATR Range Control Facility and the NASA Wallops Flight Facility site over a microwave network link

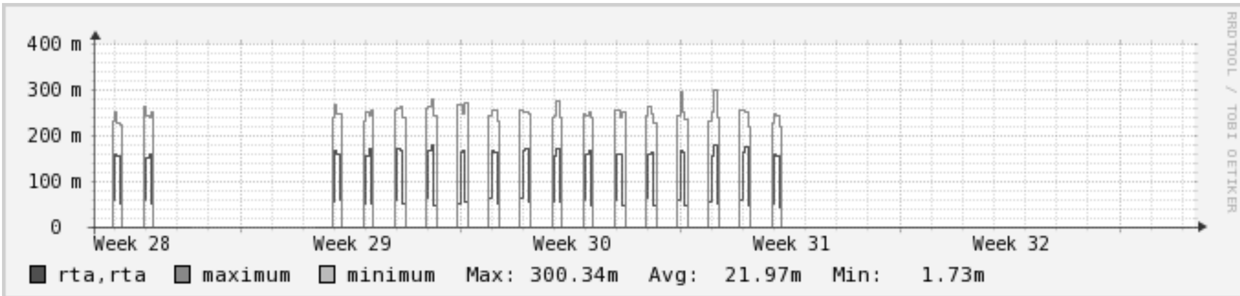

Figure 5 - Ping round trip time over the course of several weeks between the ATR Range Control Facility and the NASA Wallops Flight Facility site over a microwave network link

The fiber optic links demonstrated consistent throughput and latency throughout the test. These links experienced no downtime and can be considered highly reliable. The microwave link to NASA's Wallops Flight Facility had a somewhat variable throughput, with the link throughput degrading as much as 18%. This link's round trip time was highly variable, peaking at 4.53 ms during periods of significant delay. This is explored further in the analysis section. While this link did become degraded it never went down, and so can be considered somewhat reliable but susceptible to degradation.

All Policy tests were completed successfully. The ARON makes no changes to the DSCP set in an IP packet's header, and passes all packets regardless of what the DSCP is set to. The ARON also fully supports IGMP multicast. Both the TMS and RTNs were able to join, send data to, and receive data from multicast groups.

ANALYSIS

The data gathered in this study indicates that the fiber optic network links that comprise the majority of the ARON will easily support the implementation of several TmNSs. These links are both high capacity and high reliability. Additionally, there do not seem to be any hurdles with regards to ARON policies and the implementation of a TmNS. All the required technologies that were tested worked as needed. The microwave link from Pax River NAS to NASA's Wallops Flight Facility will support some early TmNS deployment and testing, but will not provide the same throughput or reliability as the fiber optic links.

The microwave link is an Optical Carrier-3 (OC-3) connection. While the name includes the word "optical" it is only in reference to the modulation scheme, which was originally developed for optical links. OC link bandwidth is defined according to equation 1:

$$B_{OC}("OC-N_{OC}") = 51.84 Mbps * N_{OC} \qquad (1)$$

where B_{OC} is the optical carrier bandwidth and N_{OC} is the OC number given, in this case 3. This gives a maximum link bandwidth of 155.52 Mbps. 6.912 Mbps of this are lost to overhead, leaving 148.608 Mbps of payload bandwidth. In the case of the microwave link, this is segmented into several sub-links, each for a given type of data (Video, Voice, Etc.). The Ethernet sub-link is an OC-1 connection, with a payload bandwidth of 50.112 Mbps. This differs from the observed maximum sustainable throughput by between approximately 6.5 and 13.9 Mbps. This difference is due to the nature of the Transmission Control Protocol's (TCP) flow control algorithm that was used by both the custom traffic generation tool and the iNET standards which define the TmNS. (Note: both TCP and User Datagram Protocol or UDP data were tested in this study, but TCP was studied in more depth because it is less efficient than UDP and was considered to be a worst case scenario.) TCP throughput is defined according to equation 2:

$$T_{TCP} \leq \frac{L_{Rwin}}{T_{RT}} \qquad (2)$$

where T_{TCP} is the throughput of a TCP connection, L_{Rwin} is the length of the TCP connection's Receive Window and T_{RT} is the connection's round trip time. For more on this phenomenon, see P. Dykstra's paper "*Issues Impacting Gigabit Networks: Why don't most users experience high data rates?"* [1].

A connection's receive window's behavior and size is implementation specific, so it would go against iNET's goal of vendor interoperability to try and alleviate this phenomenon by tweaking the receive window of a connection. Instead, the link should be upgraded to one that can easily host a TmNS implementation regardless of latency.

RELIABILITY

On several occasions, degradation in the microwave link was observed during weather events such as thunderstorms (Figure 6).

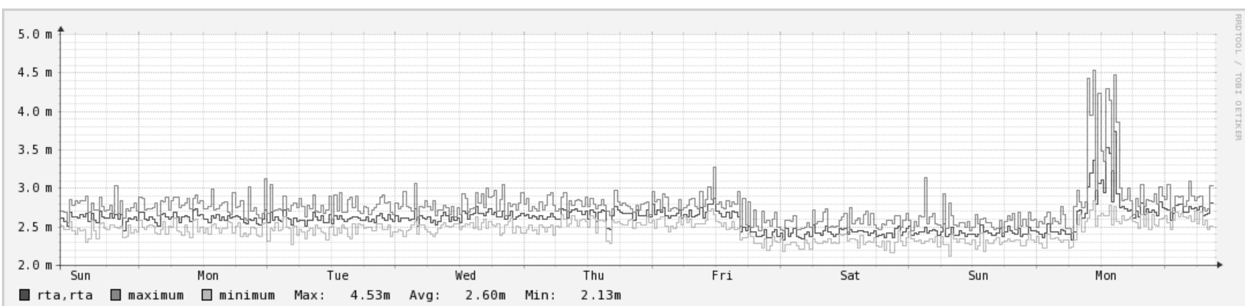

Figure 6 – Ping round trip time over the course of one week between the ATR Range Control Facility and the NASA Wallops Flight Facility site over a microwave network link showing high delay

This is due to microwave transmission's susceptibility to atmospheric interference at higher frequencies. The Wallops Flight Facility microwave link uses a carrier frequency of 23 GHz. In 2004, a study was conducted at Pax River NAS among other locations by the Jet Propulsion Laboratory and the California Institute of Technology for the Advanced Range Telemetry group at Edwards Air Force Base [2]. This study estimates the total atmospheric attenuation of a 24 GHz signal at a distance of 50 miles to be 34.3 dB, compared with 8.3 dB for a similar signal at 12 GHz. While the microwave link was never dropped completely, it did become degraded and performed at a lower level. This could be mitigated by moving to a lower frequency (an improvement of 26 dB if moved to 12GHz) in the future. This could also be mitigated by moving from a microwave link to a fiber optic link.

CONCLUSIONS

In this study, we explored the ARON, and gathered data regarding its Throughput, Reliability, and various Network Policies. From this data, we were able to construct an accurate picture of the ARON and what level of performance it was capable of. This picture has proven useful in the iNET program's continuing efforts to develop the standards which define the TmNS.

From the data collected in this study, we can conclude that the ARON is for the most part a high capacity, high reliability network which will be capable of supporting the iNET program's efforts in the immediate future. All technologies necessary for a TmNS that were tested were found to be supported by ARON policies. It was found that some antenna sites are less suited to supporting a TmNS than others. These sites should be avoided when planning future TmNS implementations until they can be upgraded.

To move forward and support future TmNS implementations, it is recommended that the microwave link between Pax River NAS and NASA's Wallops Flight Facility be upgraded to the same fiber optic OC-192 infrastructure that the rest of ARON uses. Fortunately, these upgrades are already planned as a part of the Chesapeake Bay Fiber Project. Beyond this, it is also recommended that all ranges that plan on implementing a TmNS system in the future work with iNET program personnel to develop an Upgrade Roadmap as a part of their system deployment.

It is also recommended that these tests be repeated with other ranges to develop a more complete picture of the future operating environment of the TmNS. One range particularly suited to this purpose is Pax River NAS's counterpart in the Air Force, Edwards Air Force Base. Other possible locations for future testing would include Naval Air Weapons Station China Lake, Naval Air Station Point Mugu and the Army Aberdeen Test Center.

ACKNOWLEDGEMENTS

The author would like to thank everyone who helped in this effort, including the communications staff at the Atlantic Test Range and the iNET Program Office. The author would also like to thank all the developers who contributed to the Open Source projects used in this study, especially the Fedora Project, the Nagios Network Monitoring System, and Emcast.

REFERENCES

1. Dykstra, Phil, "Issues Impacting Gigabit Networks: Why don't most users experience high data rates?", WareOnEarth Communications, Inc., San Diego, CA, Aug. 2000
2. C. Ho, C. Wang, K. Angkasa, K. Gritton, "Estimation of Microwave Power Margin Losses Due to Earth's Atmosphere and Weather in the Frequency Range of 3-30 GHz", JPL D-27879, Jet Propulsion Lab, California Institute of Technology, Pasadena, CA, Jan. 2004, Retrieved from: http://descanso.jpl.nasa.gov/Propagation/Ka_Band/JPL_D27879.pdf.