NATIONAL COMMUNICATIONS SYSTEM

Route Diversity Project (RDP) AT&T Wireless Broadband Service Evaluation Results Report

DRAFT

February 2007

Contract Number: DCA100-01-D-4001
CDRL: A032-DI-MISC-80711A

Technology and Programs Division (N2)
PO Box 4502
Arlington, VA 22204-4502

Prepared by:
Booz | Allen | Hamilton

Disclaimer:

Certain commercial equipment, instruments, or materials are identified in this document. Such identification does not imply recommendation or endorsement by the National Communications System, nor does it imply that the products identified are necessarily the best available for the purpose.

EXECUTIVE SUMMARY

This document presents the results of an AT&T wireless broadband (WBB) service evaluation conducted by the National Communications System (NCS) in partnership with BellSouth (now AT&T) as part of the NCS Route Diversity Project (RDP). The RDP is an effort initiated by the NCS to evaluate and improve the resiliency of Federal agencies' communications networks by creating methodologies and assessing technologies and services.

The evaluation was conducted from December 11 – 15, 2006 in Atlanta, GA.

PURPOSE AND OBJECTIVES

The purpose of this evaluation is to assess if the WBB service is a communications solution that can be quickly deployed in an emergency response situation and if the service is a reliable alternative for creating route diversity in the local network of Federal agencies. As part of the assessment, the NCS conducted test to determine:

- Effectiveness of the WBB service in supporting voice and data services
- Speed and ease of the Customer Premises Equipments (CPE) setup
- Maximum capacity supported by the service

SERVICE BACKGROUND

AT&T's WBB service provides wireless broadband connections to users within a certain coverage area at speeds of up to 1.5 megabits per second (Mbps). The service operates in the licensed 2.3 GHz Wireless Communication Systems (WCS) frequency band and is similar to other pre-WiMAX systems. AT&T offers the service in several areas across the nation to complement digital subscriber line (DSL) deployments. In the aftermath of Hurricane Katrina in New Orleans, AT&T used this service as a part of their disaster recovery efforts to speed the availability of broadband access to subscribers who had returned to the area.

RESULTS AND CONCLUSIONS

Based on testing results, the NCS believes that WBB service is a viable solution for both increasing route diversity and providing emergency response services within an AT&T WBB coverage area. The setup and configuration of the CPE is easy and straight forward. The link can support up to 30 users using e-mail, Web services, and FTP simultaneously. These factors make it ideal for small field offices or a small single building.

The use of WBB does have limitations, one of which is that an agency's building or remote site must be within the 3-5 mile footprint of the WBB base station. Another limitation is that while a single modem delivers satisfactory performance for a single real-time session, multiple sessions of real-time applications, such as VoIP and real-time conferencing experience degraded performance levels. Therefore, if multiple sessions are needed during emergency situations, multiple modems can be deployed.

TABLE OF CONTENTS

TABLE OF FIGURES

TABLE OF TABLES

1. INTRODUCTION

This report documents the results of an evaluation of AT&T's wireless broadband (WBB) service conducted by the National Communications System (NCS). The evaluation was conducted as part of the NCS Route Diversity Project (RDP) in partnership with BellSouth (now a part of AT&T as of January 3, 2007).

1.1 PURPOSE

The purpose of this evaluation is twofold:

- To assess the WBB service as a viable solution for enhancing route diversity in Federal agencies.
- To evaluate the WBB service's ability to be quickly deployed as an emergency response solution for providing support in the aftermath of a natural disaster event (like those relating to Emergency Service Function #2 [ESF#2]) or during a national security/emergency preparedness (NS/EP) event.

The NCS has previously evaluated other broadband wireless technologies, such as free space optics (FSO) and satellite communications (SATCOM), to determine the ability of those technologies to serve the needs of Federal agencies.

1.2 OBJECTIVES

The objectives of this evaluation are as follows:

- Observe how quickly and easily the customer premise equipment (CPE) can be setup by the user.
- Determine how effectively the WBB service can support various voice and data services.
- Verify the maximum capacity supported by the service.

1.3 SERVICE INFORMATION

AT&T's WBB service provides wireless broadband connections to users within a certain coverage area at speeds of up to 1.5 megabits per second (Mbps). The service operates in the licensed 2.3 GHz Wireless Communication Systems (WCS) frequency band and is advertised as being able to transmit signals up to 5 miles. The NCS evaluated a current WBB deployment being used by AT&T for evaluation and testing in their laboratory. The test data link is asymmetric and configured for 2 Mbps downstream and 1 Mbps upstream, higher than commercially available rates. This service also uses a proprietary wireless security scheme and supports encryption keying.

AT&T offers this service in several markets across the nation to complement digital subscriber line (DSL) deployments. In the aftermath of Hurricane Katrina in New Orleans, AT&T used this technology as a part of their disaster recovery efforts to speed the availability of broadband access to subscribers who had returned to the area. The service is currently deployed by AT&T

in many areas prone to natural disasters, such as the Southeast region of the country. However, WBB technology is not proprietary to AT&T and similar service offerings are available from other service providers. The service can be considered a predecessor to the emerging technology known as Worldwide Interoperability for Microwave Access (WiMAX).

1.4 Background Information

1.4.1 NCS Overview

As stipulated in Executive Order 12472, the NCS is responsible for the survivability and interoperability of the infrastructure supporting NS/EP telecommunications. The NCS fosters cooperation among the 23 government agencies that constitute the NCS member organizations to support their NS/EP telecommunications needs. The NCS also serves as a focal point for joint industry and government NS/EP and ESF #2 telecommunications planning and support. Additional information is available at www.ncs.gov.

The Technology and Programs Division (N2) of the NCS is responsible for administering programs, conducting technical studies, and participating in research efforts to determine how technology can better support NS/EP activities.

1.4.2 Route Diversity Project Background

The RDP is an effort by the Technology and Programs Division of the NCS to evaluate and improve the resiliency of Federal agencies' communications networks by assessing technologies and services that may provide increased resiliency. The RDP also creates tools and methodologies for increasing the communications resiliency of NCS member agencies. Additional details about the RDP can be found in Appendix H.

1.4.3 NCS-AT&T Partnership

The NCS partnered with AT&T for this service evaluation for several reasons. First, AT&T provides WBB service in places throughout the country, including several hurricane-prone areas in the Southern Atlantic seaboard and the Gulf of Mexico. In the case of inclement weather or flooding causing damage or destroying wireline infrastructure, WBB might be able to provide the connectivity needed by emergency response services. Second, AT&T is a leading provider of this service and had a previously-deployed system set up for such evaluations in its lab in Atlanta, GA. Third, both parties benefited from this partnership. The NCS was able to evaluate a potential new communication solution with minimal set-up requirements, and AT&T was provided the opportunity to have its service independently validated and to learn more about Federal communications requirements.

2. EVALUATION DETAILS

2.1 Evaluation Configuration

The NCS evaluated the wireless service by using a test deployment to which AT&T provided access. Connections to the system were provided through a laboratory located in the AT&T facility at 723 West Peachtree St. in Atlanta, GA (Site 1). The antenna and base station of the system were located on top of another AT&T building at 748 Peachtree St., a couple of blocks away from the test facility. The CPE was placed at 659 Peachtree Street (Site 2). The NCS selected the location so that the CPE was in an optimal position to receive a signal from the antenna and base station located less than a mile away. A map of the locations is shown in Figure 1.

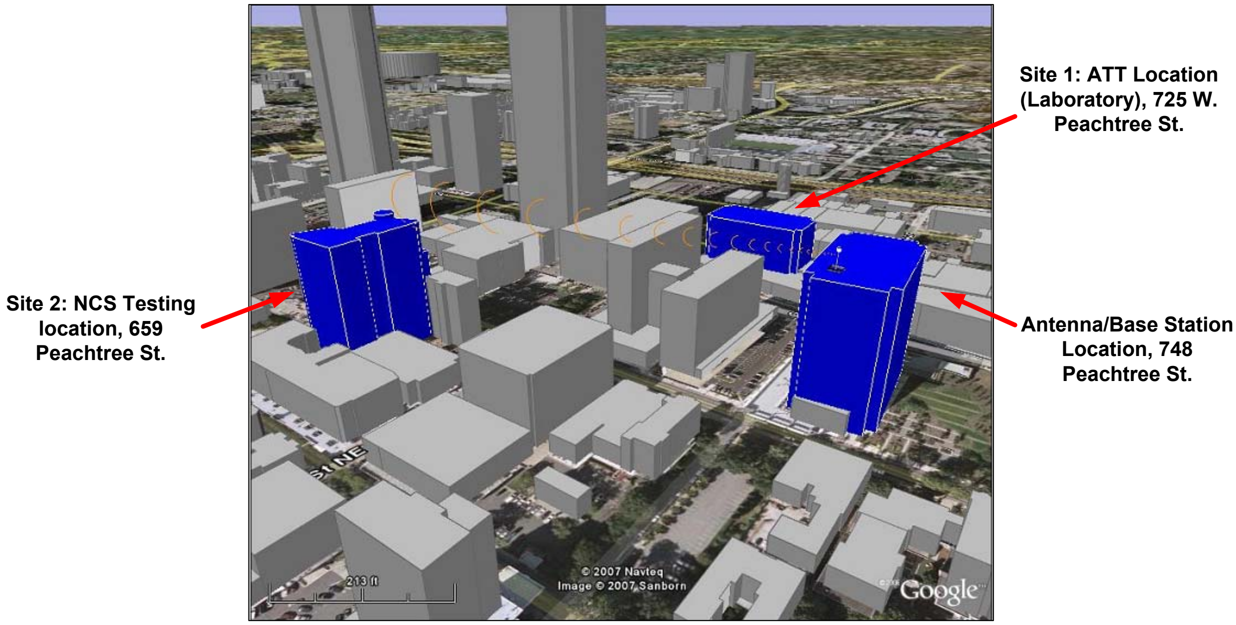

Figure 1 – Map of WBB Test Configuration

The CPE, an IXIA chassis, and a test laptop were also located at Site 2 and interconnected through a Cisco switch, as shown in Figure 2. A second IXIA chassis was placed at Site 1, where AT&T provided connections to the base station. The IXIA chassis at Site 1 was connected to the base station through an Ethernet hub provided by AT&T.

The test laptop was used to control and setup tests on both IXIA chassis, as well as to collect the test results. The IXIA chassis generated and received test traffic at both locations across the WBB link. Each test chassis was managed using Virtual Network Computing (VNC) remote desktop software, which allowed for remote configuration without multiple monitors and keyboards. All tests were performed in “batch mode”, which allows the Ixia equipment to use the same WBB link for both transmission of test data and collection of test results/management traffic between the two Ixia chassis. The “batch mode” ensured no management traffic was on the link during testing. The test configuration is depicted below in Figure 2 and a detailed description of the test equipment and software is in Appendix D.

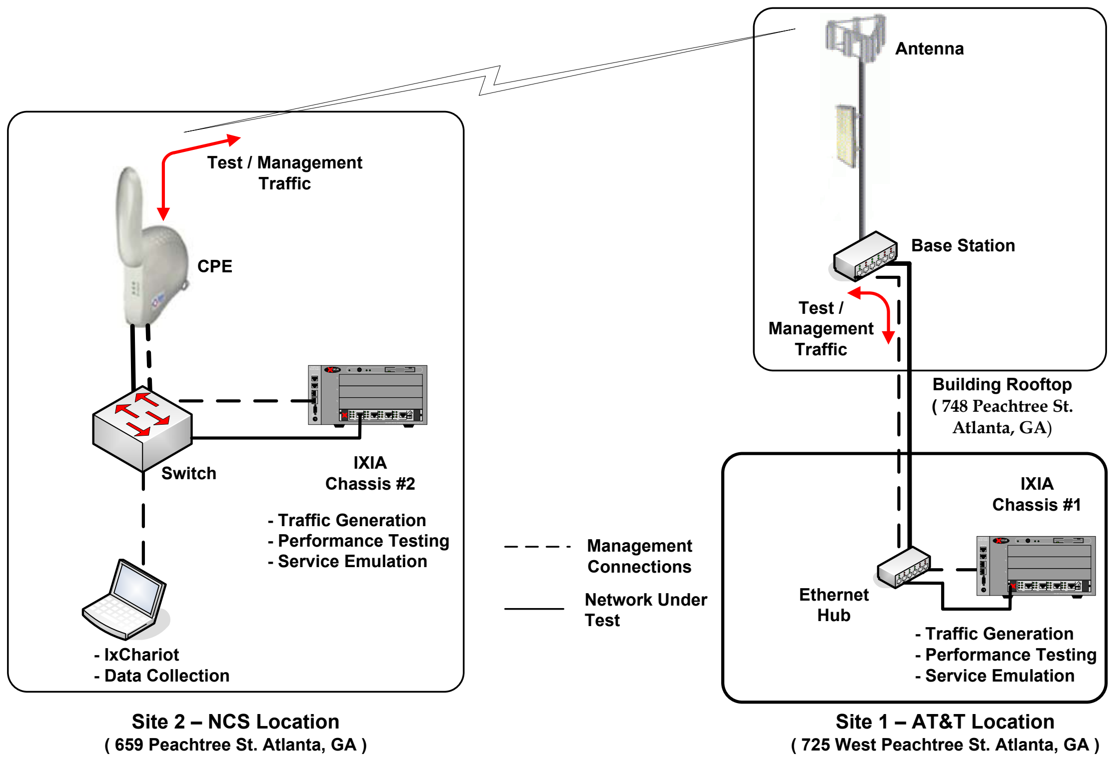

Figure 2 – WBB Test Configuration

2.2 Tests Performed

The NCS evaluated AT&T's WBB by conducting performance testing and services emulation testing. This section describes the tests performed. For greater detail about the tests performed and the results of those tests, see Appendices B and C.

2.2.1 Setup Observations

During the CPE setup, observations on the ease of setup and configuration of the system were noted.

2.2.2 Performance Testing

Performance testing measured the throughput of the WBB link. Two types of throughput testing were conducted.

RFC 2544 Throughput Test

The first test used IxExplorer to benchmark the WBB link throughput in accordance with the RFC 2544[1] specification at Layers 1 - 3 of the protocol stack. The RFC 2544 tests provide

[1] http://www.rfc-archive.org/getrfc.php?rfc=2544

network equipment users with a set of standardized, vendor-independent tests for use in comparing equipment from different vendors.

IxChariot Throughput Test

In the second test, IxChariot was used to determine the throughput performance of applications using Transmission Control Protocol/Internet Protocol (TCP/IP). IxChariot simulates applications that use protocols at all layers of the protocol stack, providing a more comprehensive evaluation including impact on performance due to TCP features such as error detection, windowing, sequencing, and retransmissions not provided by IxExplorer.

2.2.3 Services Emulation Testing

The purpose of the service emulation testing was to provide an assessment of the transmission quality and performance characteristics for various applications across the WBB link. It also provided a baseline for expected performance of multiple applications being used simultaneously across the wireless link. The IxChariot test tool was used to configure, emulate, and collect statistics on the following service applications:

- File Transfer Protocol (FTP)
- E-mail
- Web (HTTP)
- Multiple services (e-mail and Web services being used simultaneously)
- Voice over Internet Protocol (VoIP)
- Streaming media

2.3 Test Environment Constraints

A few test environment constraints potentially affected evaluation results. First, the AT&T WBB technology did not have quality of service (QoS) configured for real-time traffic; as a result, the technology was not optimized for VoIP and streaming applications. Second, AT&T noted that there were other subscribers using the WBB technology; therefore, NCS test traffic was not the only traffic present on the network. Traffic from other subscribers may have affected the results.

3. RESULTS

3.1 Installation

The only equipment setup required was the installation of the CPE, which was straight forward. The CPE came preconfigured with the log-on information necessary to authenticate itself with the base station. Once powered on, the CPE automatically established a connection with the base station. The system did not require a line of sight to the antenna, thus no special positioning of the CPE was necessary to receive a signal. The quick CPE setup is particularly useful for fast deployment during a disaster. The system supports encryption, but encryption was not configured or tested during this evaluation.

3.2 PERFORMANCE TESTING

The NCS evaluation of the AT&T WBB link verified that the downstream and upstream throughput rates were in line with the rates specified by AT&T. Both measurement techniques, RFC 2544 and IxChariot, provided consistent performance levels. The NCS measured slightly lower downlink speeds and slightly higher uplink speeds than rated speeds. The rated link speeds for the WBB test link were 2 Mbps downstream and 1 Mbps upstream. Measuring throughput in accordance with the RFC 2544 specification, the average throughput of all frame sizes was 1.745 Mbps (87.25% of 2 Mbps) in the downstream direction and 1.083 Mbps (108.3% of 1 Mbps) in the upstream direction.

When measuring throughput utilizing the entire protocol stack and forwarding traffic in one direction only, the average throughputs measured were 1.837 Mbps downstream (91% of 2 Mbps) and 1.176 Mbps upstream (117% of 1 Mbps). This test only used the maximum frame size of 1518 bytes when sending data. Transmitting traffic in both directions simultaneously resulted in the throughput falling significantly. The total throughput was found to be only about 40% of the total (1.2 Mbps against 3 Mbps) available bandwidth and was seemingly operating in half-duplex. An Ethernet hub to which one of the Ixia chassis was connected is the suspected cause for this behavior perhaps due to its half-duplex operation.

Table 1 - Throughput Testsing Results

Test	Direction	Rated Link Speed	Throughput	
			Average of all Frames*	Frame Size: 1518
RFC 2544 Throughput Test	Downstream	2 Mbps	1.745 Mbps	1.86 Mbps
	Upstream	1 Mbps	1.083 Mbps	1.27 Mbps
IxChariot Throughput Test	Downstream	2 Mbps	N/A	1.83 Mbps
	Upstream	1 Mbps	N/A	1.17 Mbps

* Frame Sizes 64, 128, 256, 512, 1024, 1280, 1518

3.3 SERVICES EMULATION TESTING

Multiple Services (FTP, Web, and Email)

The results of the FTP application test verified that FTP transactions could be conducted successfully across the link. The FTP throughput results also confirmed the throughput test results. The HTTP application test verified no unusual response times during the transmission of a single 530 kilobytes (k) web page with text and graphics. The average response time for the 30k text segment of the webpage was 0.824 seconds and for the 500k graphics segment it was 5.785 seconds. The email application test verified the successful transmission of large and small emails using both the SMTP and POP3 protocols.

The multiple services test replicated normal data traffic and measured the number of users that the AT&T WBB link could support. Based on the combined performance of email and web browsing applications (HTTP) for a various number of simultaneous users, a WBB link rated at 2

Mbps downstream and 1 Mbps upstream is able to support up to 30 simultaneous users at a remote site.

VoIP

Due to the way codecs convert voice into bits, they have theoretical maximum Mean Opinion Score (MOS) limits. The G.711u codec can only provide a maximum MOS of 4.4 and the G.729 codec's MOS cannot exceed a value of 4.07. Due to losses in conversion, no codecs ever provide an ideal MOS of 5.0. For this testing, a call was considered unsatisfactory if its MOS was below 3.10. The VoIP performance was satisfactory for only a single call using the G.711u codec. Adding a second VoIP call brought down the MOS to a value of 2.40. The lesser quality G.729 codec did not provide acceptable VoIP performance for a single call. One side of the call experienced a MOS value of only 2.67. The VoIP performance in each case was affected by jitter, one way delay, and data loss. AT&T has not optimized their WBB links for VoIP performance and other real-time sensitive services. This could be a possible reason for the poor VoIP performance. Configuring the technology for quality of service and support for real-time sensitive services may improve VoIP performance. Since each modem is limited in the number of real-time sessions that it can support, multiple modems can be used to support multiple VoIP sessions in emergency situations.

Real-Time Streaming

Real-time streaming applications were emulated across the WBB link using Microsoft NetMeeting. The results of the tests demonstrated that streaming media performance is satisfactory for only a single media stream. Just as with VoIP, streaming media applications are equally sensitive to jitter, delay, and data loss. The single media stream performed within the thresholds outlined in Table 2. Two simultaneous NetMeeting sessions degraded both sessions to levels beyond the tolerances outlined in Table 2. The percentage of bytes lost often exceeded 15% and the jitter and one-way delay often exceeded 60 milliseconds (ms) and 200 ms respectively. This effect can be attributed to the fact that the WBB link is not configured with a QoS and is not optimized for real-time services. As indicated in the VoIP section, real-time streaming performance degradation can be minimized by using multiple modems.

Table 2 – Real-Time Streaming Requirements

Metric	Tolerances
One-Way Delay	< = 150 ms
Jitter	< = 30 ms
% of Bytes Lost	< = 1 %

4. CONCLUSION

Based on this evaluation, the NCS believes that AT&T's WBB can provide both a route diversity solution and an emergency response solution for Federal agencies. However, the limitations discussed below should be considered for deployment of the service.

Any use of the service is confined within 3-5 mile range of an AT&T WBB base station (or the base station of a company providing similar WBB service). For a route diversity service, an agency must confer with AT&T or other WBB providers to see if it is within the coverage area. If the agency is within a coverage area, then WBB can provide a back-up communications capability for agencies in the case of wireline infrastructure failure, thereby increasing communications resiliency at an agency facility.

A single WBB modem easily supports a single real-time session such as VoIP or streaming. However, the performance of the service on a single modem is degraded for multiple sessions of real-time applications, such as VoIP and real-time conferencing. Therefore, multiple modems with differentiated channels can be deployed to meet capacity demands of the multiple sessions required during emergency situations.

WBB can also provide emergency response services. It is easily deployable, so long as the CPE is in the 3-5 mile footprint of the base station, requiring little more set-up than plugging in the CPE and connecting it to computers. Agencies looking to deploy such services in emergency situations should ensure the system is configured and optimized for real-time applications.

APPENDIX A - ACRONYMS

CRC	Cyclic Redundancy Check
CPE	Customer Premise Equipment
DSL	Digital Subscriber Line
FSO	Free Space Optics
FTP	File Transfer Protocol
GHz	Gigahertz
GUI	Graphical User Interface
HTTP	HyperText Transfer Protocol
IP	Internet Protocol
LAN	Local Area Network
Mbps	Megabits per second
MOS	Mean Opinion Score
NCS	National Communications System
NS/EP	National Security/Emergency Preparedness
QoS	Quality of Service
RDP	Route Diversity Project
RFC	Request For Comments
RTP	Real-Time Transport Protocol
SPX	Sequenced Packet Exchange
TCP	Transmission Control Protocol
UDP	User Datagram Protocol
VNC	Virtual Network Computing
VoIP	Voice over IP
WBB	Wireless Broadband
WCS	Wireless Communications System
WiMAX	Worldwide Interoperability for Microwave Access

APPENDIX B - PERFORMANCE TESTING DETAILS

RFC 2544 THROUGHPUT TEST

The throughput test used IxExplorer to benchmark the WBB link throughput in accordance with the RFC 2544[2] specification at Layers 1 - 3 of the protocol stack. The RFC 2544 tests provide network equipment users with a set of standardized, vendor-independent tests for use in comparing equipment from different vendors. The RFC 2544 specifications outline the methodology for four network benchmarking tests. Of the four tests, only the throughput test meets the goals of this evaluation. The throughput test measured the maximum transmission rate with no frame errors or loss by increasing or decreasing the line rate automatically until all frames were transmitted and received successfully. This test was performed in both directions independently to find the maximum upstream and downstream throughputs for forwarding frames. The test was completed for each of the standard Ethernet frame sizes 64, 128, 256, 512, 1024, 1280, and 1518.

Test Results

The rated bandwidth of the WBB link in the downstream direction is 2 Mbps. The RFC 2544 test was set to start with a 5 Mbps traffic load. For each frame size, the traffic load rate was incrementally reduced until a throughput threshold with no errors was reached. Figure 3 shows the maximum error-free throughput recorded for various Ethernet frame sizes in the downstream direction. The results of the test do not show any particular correlation between throughput and frame size; therefore, no formal conclusion could be made based on frame size. The average throughput of all the frames combined was about 1.573 Mbps (78.65% of 2 Mbps). By excluding the 256 byte size frame result, because its result seemed to be an anomaly, the average throughput came out to be 1.745 Mbps (87.25% of 2 Mbps).

The rated bandwidth of the WBB link in the upstream direction is 1 Mbps. As in the downstream test, to capture the maximum possible bandwidth the RFC 2544 test was set to start with a 5 Mbps traffic load. Figure 4 shows the maximum error-free throughput recorded for each of the Ethernet frame sizes in the upstream direction. With the exception of the 1280 byte size frame, the results of the test showed a trend of increasing throughput as the frame size increased. The average throughput of all the frame sizes came out to be 1.083 Mbps and larger than the rated 1 Mbps speed (108.3% of 1 Mbps).

As mentioned in the test environment constraints section, the AT&T base station was shared by a number of subscribers. Some of the throughput performance fluctuations may be attributed to traffic demands from other subscribers or interference on the link.

[2] http://www.rfc-archive.org/getrfc.php?rfc=2544

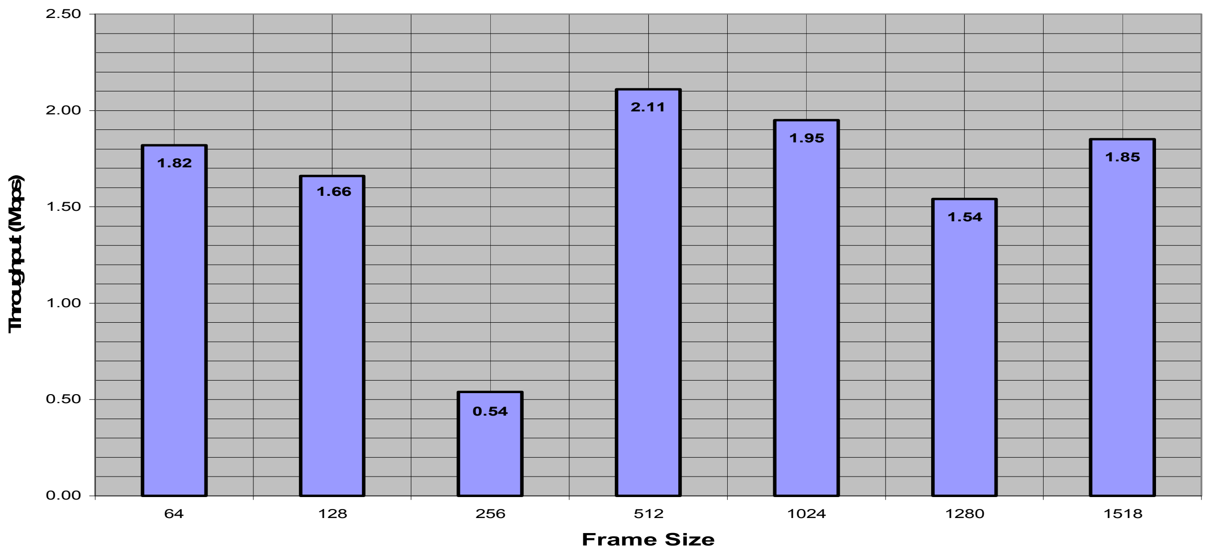

Figure 3 - RFC 2544 Throughput (Downstream)

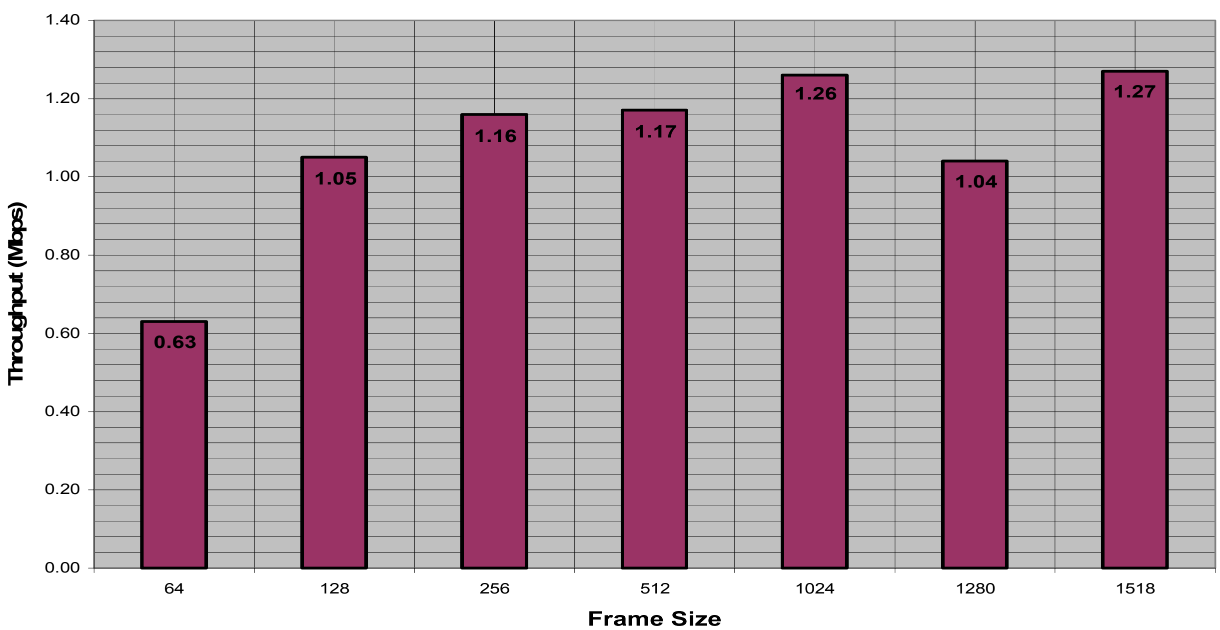

Figure 4 - RFC 2544 Throughput (Upstream)

IXCHARIOT THROUGHPUT TEST

Once the throughput of the lower layers was identified, the WBB link's throughput performance was also evaluated by an IxChariot test. The IxChariot throughput test was used to determine the throughput performance that a typical application using TCP/IP will experience. The RFC 2544 performance test only sends generic IP layer datagrams (Layer 3), while IxChariot simulates services that use protocols at all layers of the protocol stack. It also provides a more comprehensive evaluation, including impact on performance due to TCPs features such as error detection, windowing, sequencing, and retransmissions not provided by IxExplorer. The IxChariot test measures throughput using only the maximum frame size of 1518 bytes. This test was carried out in two parts. The first part consisted of sending a 10 MB file in each direction at different times without dictating any other parameters (such as frame sizes), five consecutive times. In the second part, 10 MB files were sent and received simultaneously, five consecutive times – again without dictating any other parameters.

Test Results

The IxChariot throughput results demonstrated that the link performed as specified by AT&T for both the upstream and downstream single directional tests. The technology was rated at 2 Mbps downstream and 1 Mbps upstream and the results showed an average of 1.837 Mbps downstream and 1.176 Mbps upstream. When traffic was transferred in both directions simultaneously (full-duplex), the upstream and downstream throughputs were significantly reduced. In the full-duplex test, the average downstream throughput was 0.906 Mbps (45% of 2 Mbps) and the upstream throughput was 0.177 Mbps (18% of 1 Mbps) consuming a total of only 1.08 Mbps (36% of the combined 3 Mbps link maximum speed) of the rated bandwidths combined. The WBB link is a full-duplex link; a possible cause of the half duplex behavior was an Ethernet hub in the network path; hubs restrict traffic to half-duplex operation. Figure 5, shown below, graphs the average throughput from the single direction and full-duplex tests simultaneously along with the total upstream and downstream throughputs added together.

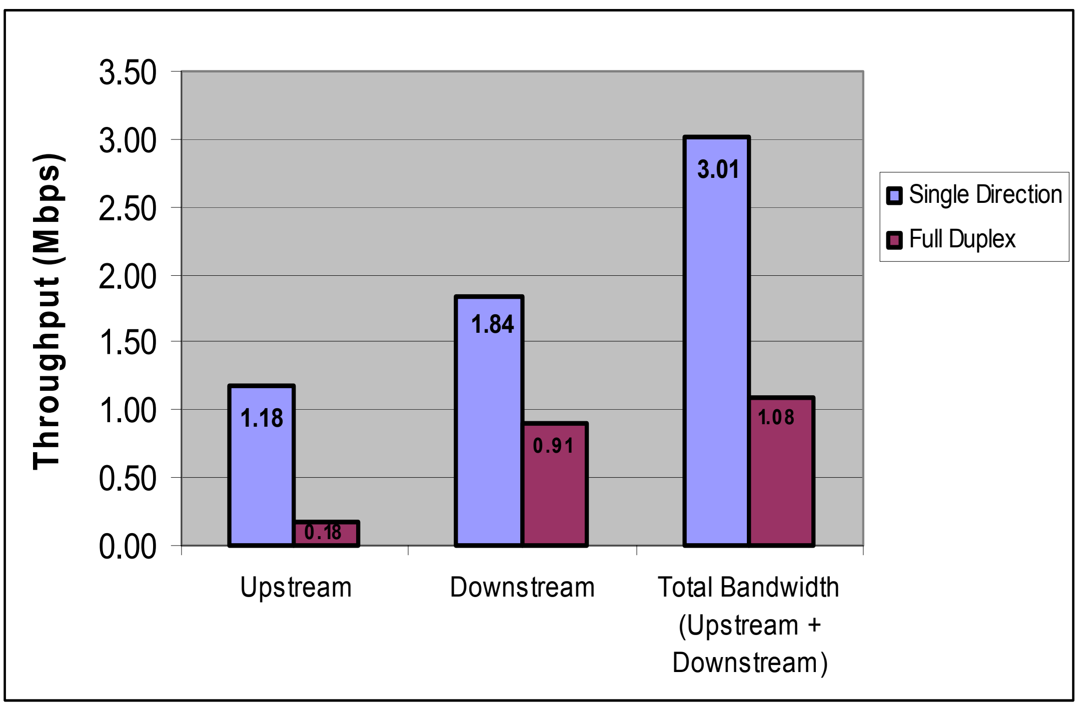

Figure 5 - IxChariot Average Throughput

APPENDIX C – SERVICES EMULATION TESTING

The purpose of service emulation testing is to provide an assessment of the transmission quality for various applications across the WBB link. It also provides a baseline for expected performance characteristics of specific applications running across the wireless link. The IxChariot test tool was used to configure and emulate various applications and collect statistics for the tests described in the sections below.

FTP

The FTP service test evaluated the performance of FTP applications on the link. IxChariot emulated FTP transactions between one FTP server and a single FTP client by emulating all required traffic flows of a FTP application. The traffic flows consisted of logging into the server, selecting a file, and transferring a 10 MB file. The FTP Test transferred the 10 MB file in the upstream and downstream direction independently, then in both directions simultaneously (full-duplex). For each run, the test performed the FTP transfers five times to get a sample of the performance over multiple runs. Table 1 shows the location of the FTP server and FTP client for each test.

Table 3 – FTP Configurations

Direction	Endpoint 1 (Site 1)	Endpoint 2 (Site 2)
Upstream	FTP Server	FTP Client
Downstream	FTP Client	FTP Server

Test Results

The results of the FTP service test validated the throughput rates already measured across the link by the previous tests. As expected, the throughput results match the results obtained in the IxChariot throughput test. Also similar to the IxChariot throughput test, full-duplex performance was limited to seemingly half-duplex operation. The average throughput results of the test when run in each direction independently and in full-duplex are shown in Figure 6.

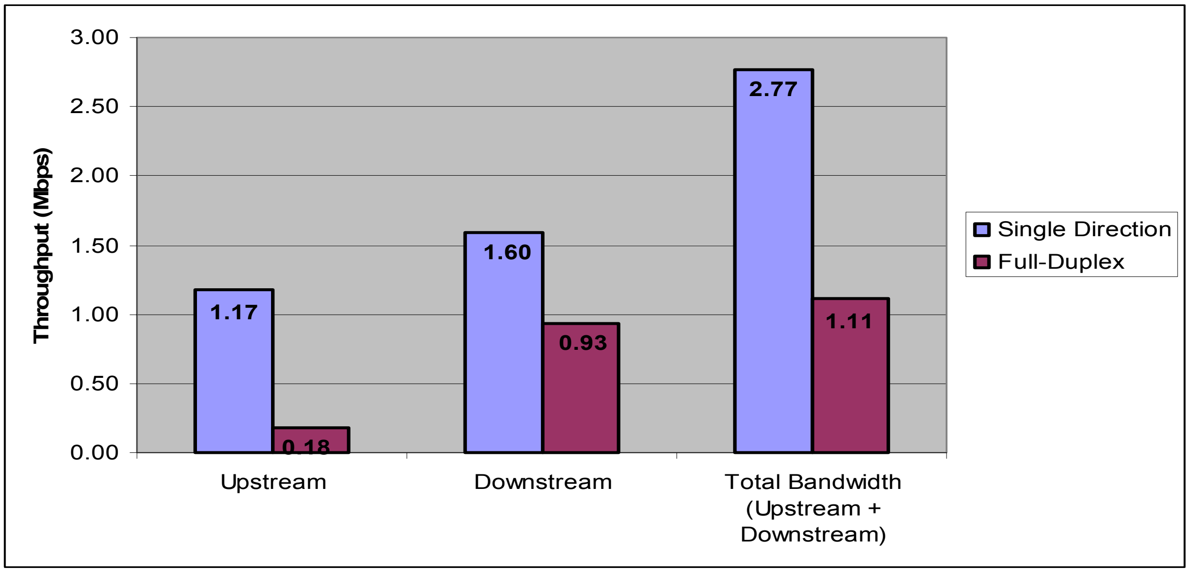

Figure 6 - FTP Average Throughput

E-MAIL

The e-mail service test evaluated how email transactions would perform across the link using two standard email transfer protocols: SMTP and POP3. The test emulated all traffic flows generated by each of the protocols for the sending and receiving of emails. In order to simulate large and small emails, two email sizes were selected: 35 kb and 1 MB.

Test Results

The results of the email services test demonstrate that email services using SMTP and POP3 can be used on the link. The response times of the email transactions for each of the file sizes in each direction are shown in Figure 7. The smaller emails of 35 kb experienced the same response time in both directions. The larger emails of 1 MB had shorter response times in the downstream direction and longer response times in the upstream direction. This behavior was expected given the differing link speeds of 2 MB downstream and 1 MB upstream.

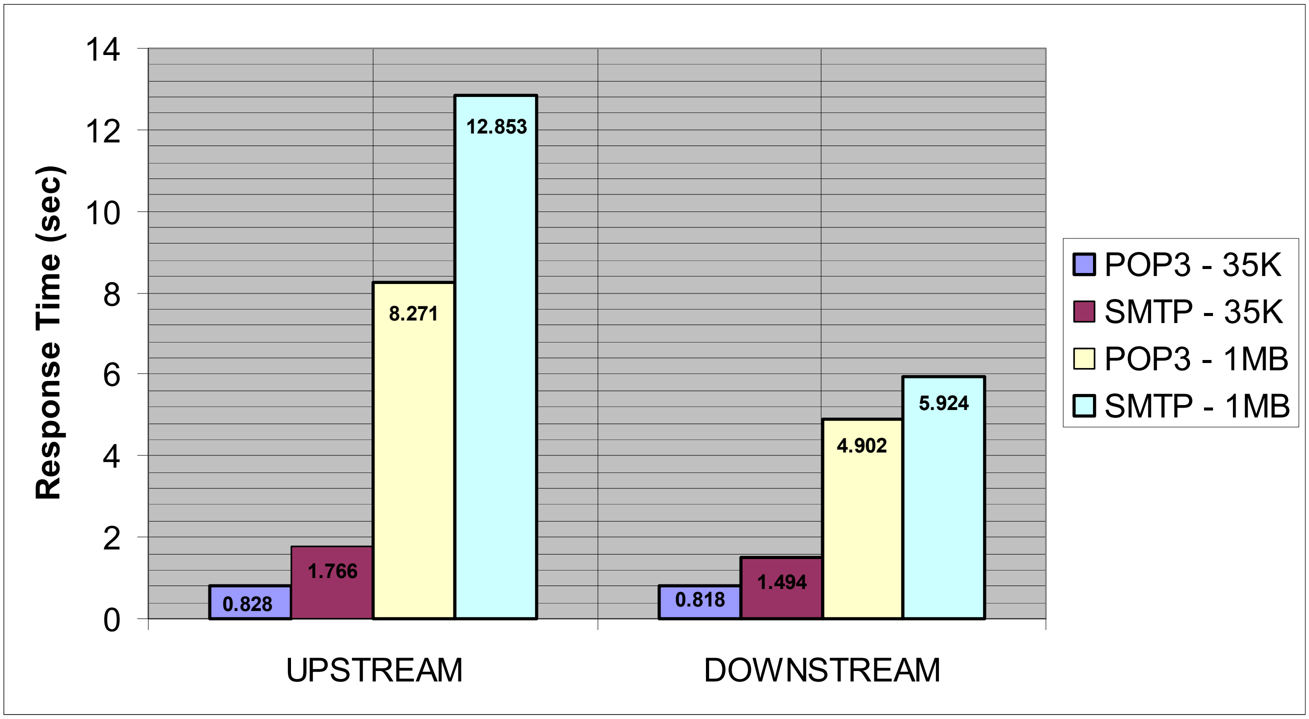

Figure 7 – Email Average Response time

WEB (HTTP)

The HTTP Test emulated transactions between a HTTP server on one side of the link (Site 1) and a HTTP client on the other side of the link (Site 2). The test emulated the transfer of web pages with text and graphics. Transfer of text and graphic files were done independent of each other. For both text and graphics test emulations, the test was run 50 consecutive times to get large sample of data.

Test Results

The results of the HTTP test verified that no unusual response times were experienced by the users of this service. The average response time for the 30k text segment of the webpage was .824 seconds and for the 500k graphics segment it was 5.785 seconds. Figure 8 displays the average response time for the HTTP transactions.

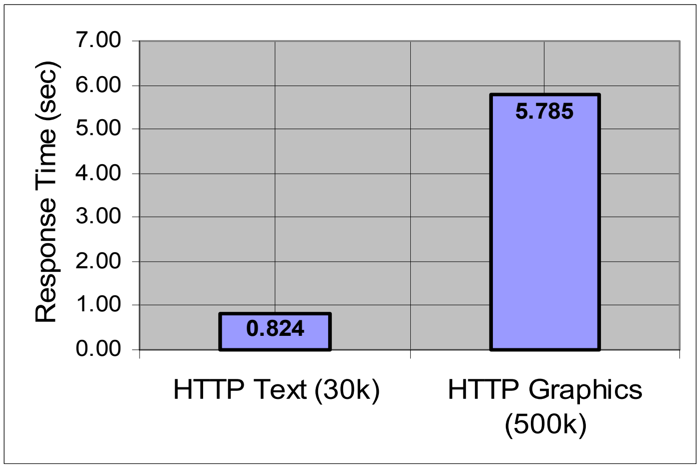

Figure 8 – HTTP - Average Response Time

MULTIPLE SERVICE TESTING

The multiple services test emulated users performing the two common data transactions of using email and web browsing (HTTP) in a typical NCS scenario. In a typical scenario, users are connected to the CPE behind a switch with the majority of the traffic composed of users receiving HTTP traffic downstream (in the user direction), users receiving email traffic in the downstream direction, and users sending outbound email traffic in the upstream directions. The test began by emulating only one user using HTTP and email services simultaneously. The test then increased the number of concurrent users sending traffic according to this profile to 10, then 20, and concluded with 40 users. The test used a normal distribution to emulate users randomly sending and receiving emails and browsing web pages. This was done to make this traffic profile more realistic.

Test Results

The results of the multiple services test are summarized in Figure 9. The response time remains under an acceptable 11 seconds for each of the services up to 20 users. In each case the downloading of the graphics webpage had the largest response time which was expected due to its large size of 500k. The results also show that the response times for each of the services deteriorated greatly between 20 and 40 users.

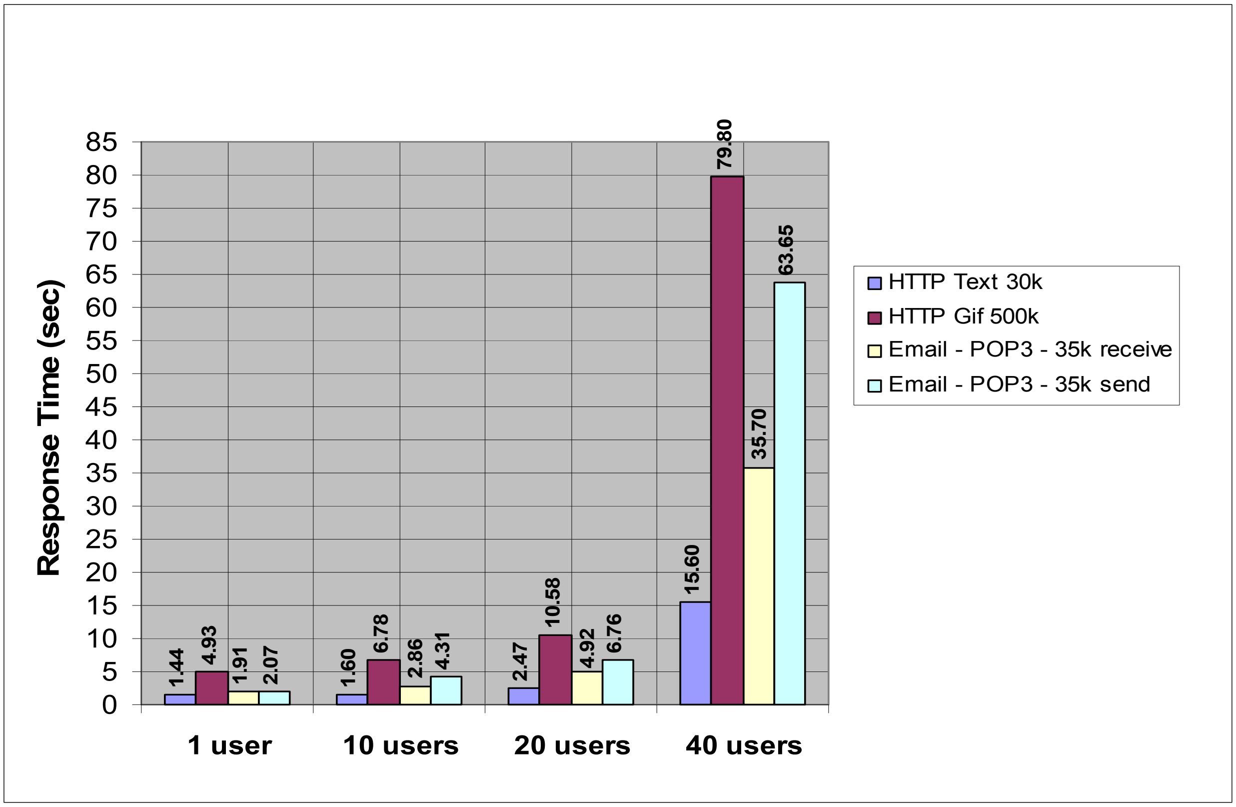

Figure 9 – Multiple Services Test - Avg Response Time

VOIP TESTING

The VoIP test was used to determine the maximum number of full-duplex and uni-directional VoIP calls that could be supported by the link while maintaining good call quality levels for the voice. Two VoIP codecs (G.729 and G.711u) were selected for this evaluation. For each codec, the test measured the quality of a single voice call. After the initial single call, additional calls were added until the quality of the voice calls became unsatisfactory. The performance of the G.729 and G.711u codecs was also compared. The overall VoIP call quality was rated using the MOS rating system and calls were characterized as unsatisfactory when the MOS fell below 3.10. IXIA's IxChariot estimates MOS scores by using factors such as jitter, latency, and data loss to determine a score that predicts user satisfaction. Each voice call simulated speech between 5 and 7 seconds and was repeated 100 - 200 times to give a large sample of data. The VoIP performance was tested without any other traffic on the link. Throughout these results, the terms upstream and downstream correspond to the 1 Mbps and 2 Mbps link bandwidths respectively.

Test Results

711u Codec – 1 call

The VoIP test emulating a single full-duplex call using the 711u Codec demonstrated that one full-duplex call can be placed on the link with relatively decent quality. Ixia measured the voice quality in upstream and downstream directions independently. The upstream call had an average MOS score of 3.39 while the downstream call score was 4.18 (Figure 10).

The two factors that affected the voice quality most frequently were jitter and one way delay. Jitter is the variation in the one way delay. To correct jitter, VoIP devices use a jitter buffer to briefly store packets and reorder them if necessary. The jitter buffer then releases the packets at evenly spaced intervals, thus removing delay variation. Packets with excessive delay due to jitter are dropped by the buffer and are not received by the receiver. The jitter buffer size used for this test was 40 ms. The evaluation results show that during the periods when the number of packets dropped by the jitter buffer increased, the MOS value decreased. Packets are transmitted in the order of milliseconds, thus there are many packets sent per second. To measure the jitter of these packets, the RFC 1889 jitter measurement was used. The RFC 1889 jitter measurement calculates a mean statistical deviance of the packets inter-arrival times (jitter) and is shown in Figure 12. Given that the RFC 1889 jitter shows an average level, the jitter seems acceptable and it is not obvious why the jitter buffer drops packets. The maximum jitter measurements shown in Figure 13, gives better insight by showing that many of the packets had jitter over 50 ms. The one-way delay time in many cases also exceeded the recommended maximum threshold for VoIP of 150 ms as shown in Figure 11.

The evaluation results also demonstrated that the upstream link had higher jitter and forwarded datagrams with higher one way delay. Overall, more datagrams and bytes were lost in the upstream direction. The upstream link had a lower MOS rating than the downstream portion of the call. Table 4 summarises the upstream and downstream MOS values for this test.

The throughput for the calls remained relatively fixed and only showed strong dips when there was noticeable data loss. Figure 15 provides a graphical view of the data loss and Figure 16 graphs the throughput of the link. The data loss graphed in Figure 15 represents data that is lost before it has reached any of the second endpoint's buffers (e.g. jitter buffer). These graphs highlight two periods of data loss correlated with dips in the throughput at the beginning and at approximately 12.58 minutes into the test. During these two periods of data loss, the MOS also dropped to a value of approximately equal to 1. These dips had very short durations and their causes are unknown.

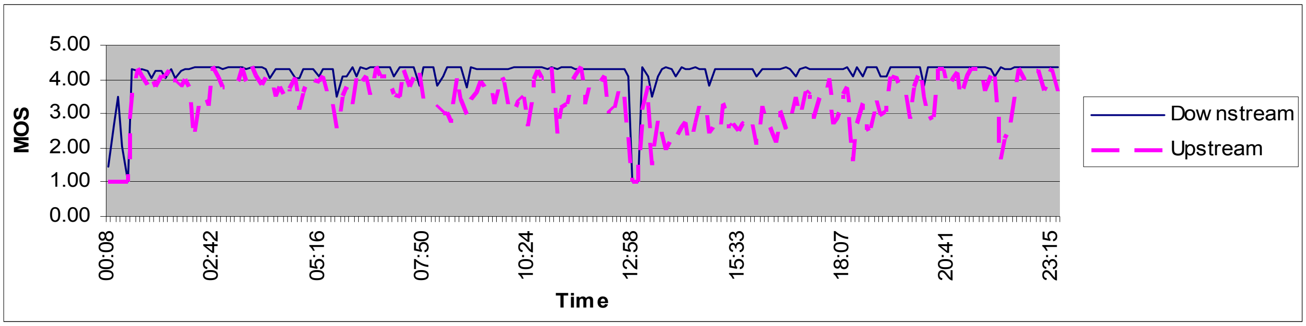

Figure 10 - MOS (1 call – G711u)

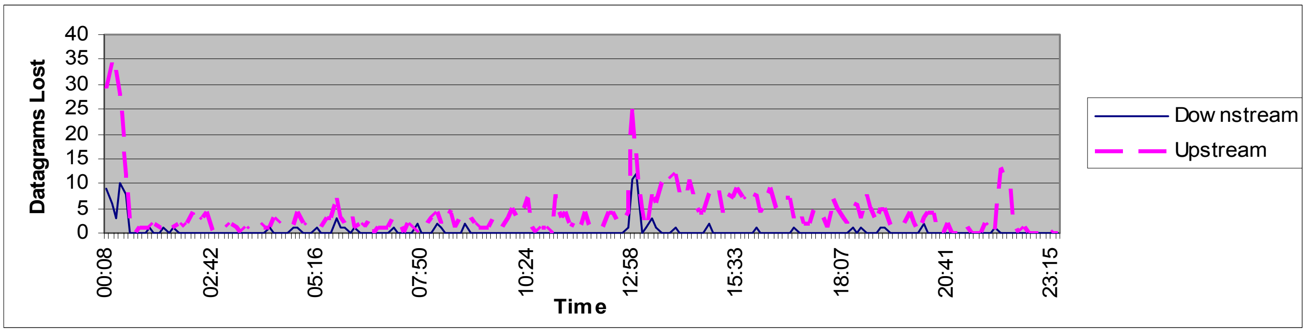

Figure 11 - Jitter Buffer Datagrams Lost (1 call – G711u)

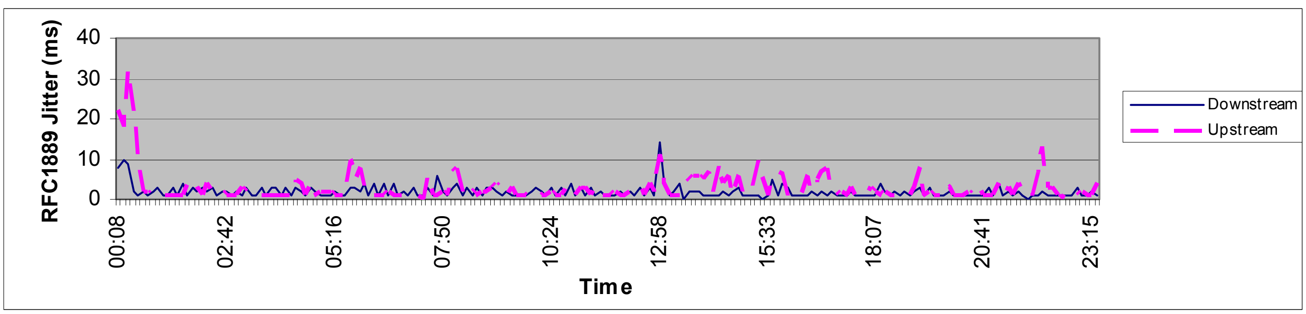

Figure 12 - RFC 1889 Jitter (1 call – G711u)

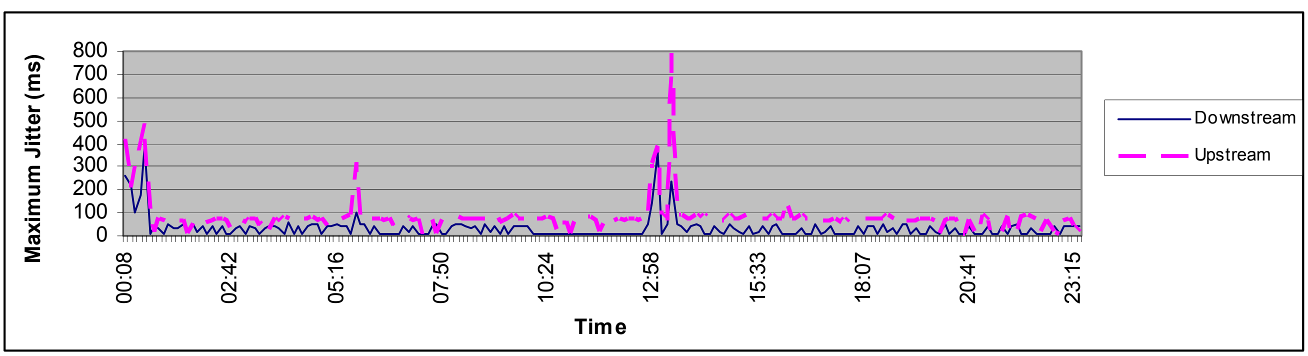

Figure 13 – Maximum Jitter (1 call – G711u)

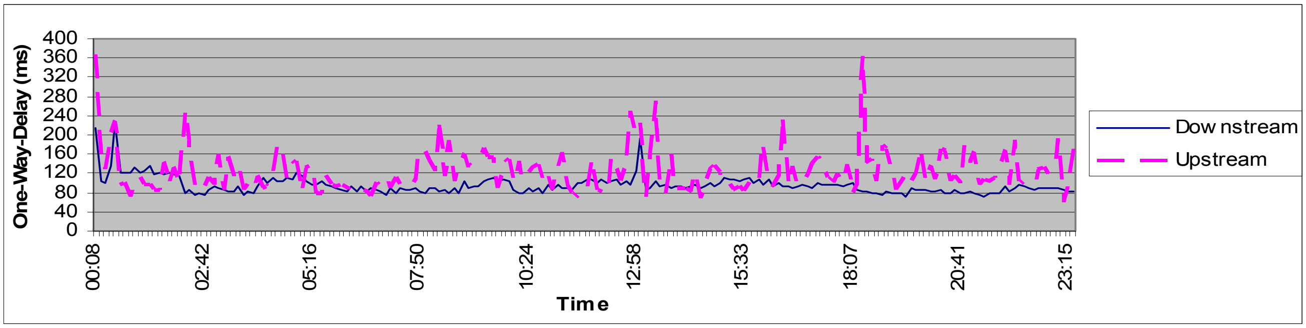

Figure 14 - One Way-Delay (1 call – G711u)

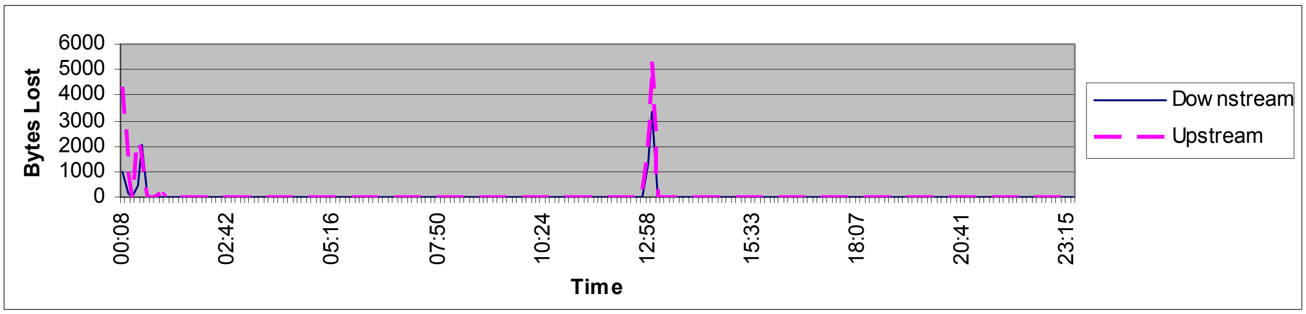

Figure 15 - Bytes Lost (1 call – G711u)

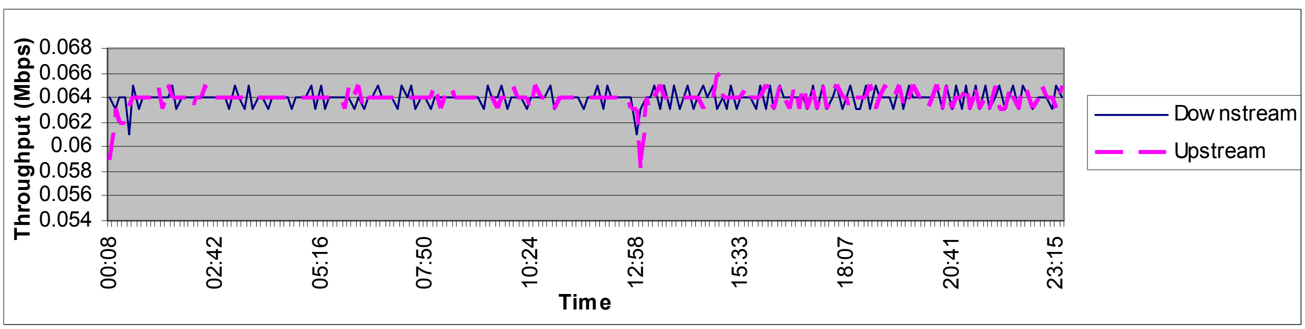

Figure 16 - Throughput (1 call – G711u)

711u Codec – 2 calls

In the next portion of the evaluation, an additional VoIP call using the 711 Codec was added to measure if the link could support a second VoIP call. The VoIP calls in the two call testing scenario experienced a lower MOS and provided noticeably lower call quality. The differences in MOS values from the two call test versus the single call test are shown in Figures 17 and 20. The calls in the two-call test experienced greater jitter and one way delay values than in the single call test. The calls in this test also experienced more frequent and longer durations of significant data loss than in the previous single call test. Figure 17 through Figure 30 provide a side-by-side presentation of the statistics of the single call test versus one of the calls from this two-call test. The MOS results are summarized in Table 4.

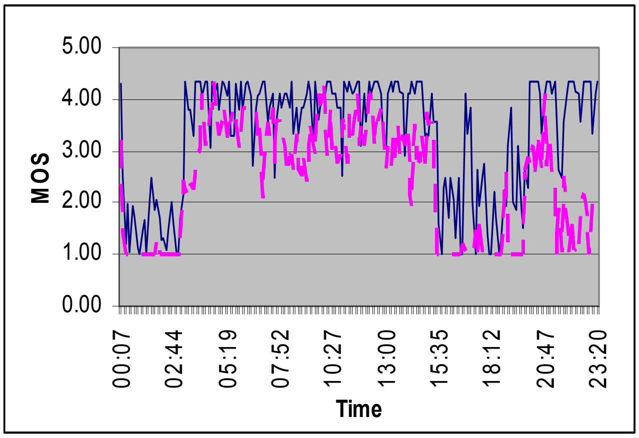

Figure 17 - MOS G711u (1st of 2 Calls)

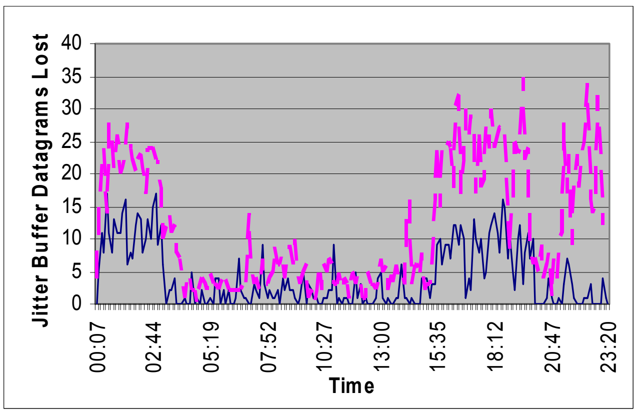

Figure 18 - Jitter Buffer Datagrams Lost G711u (1st of 2 Calls)

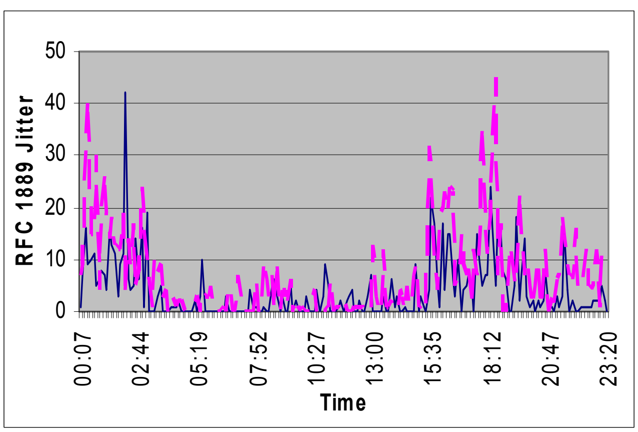

Figure 19 - RFC 1889 Jitter G711u (1st of 2 Calls)

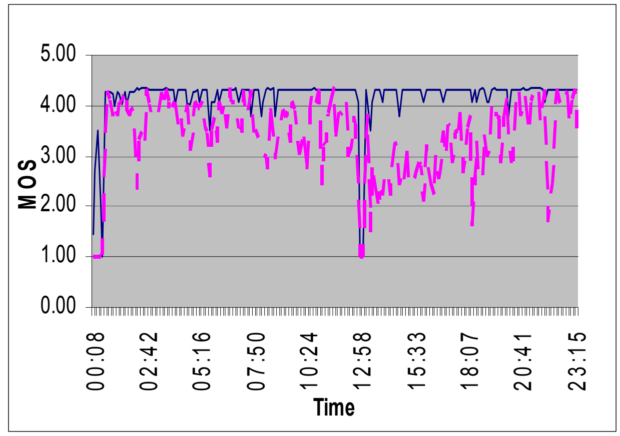

Figure 20 – MOS G711u (Single Call)

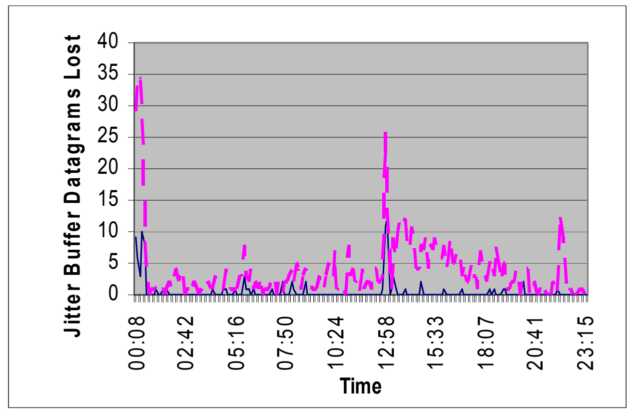

Figure 21 - Jitter Buffer Datagrams Lost G711u (Single Call)

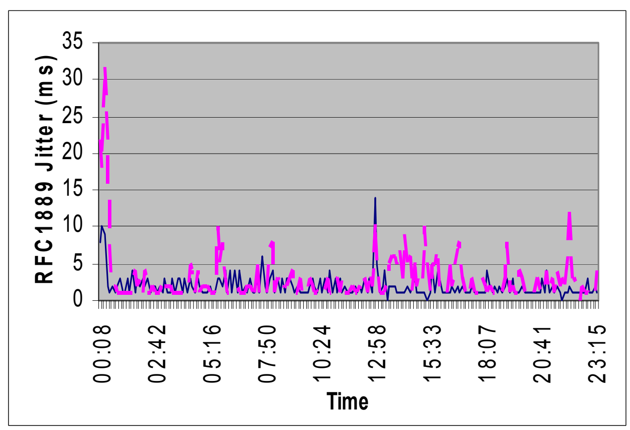

Figure 22 – RFC 1889 Jitter G711u (Single Call)

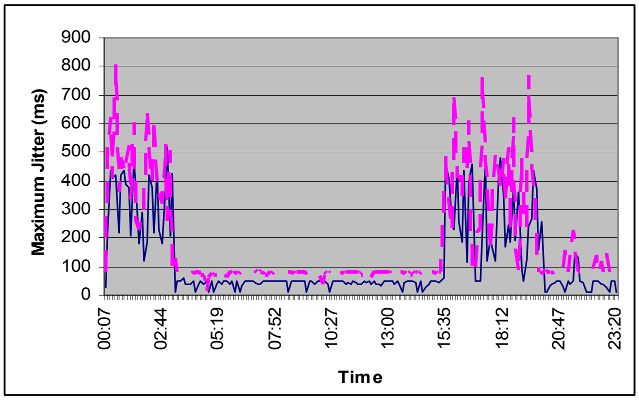

Figure 23- Maximum Jitter (1st of 2 Calls)

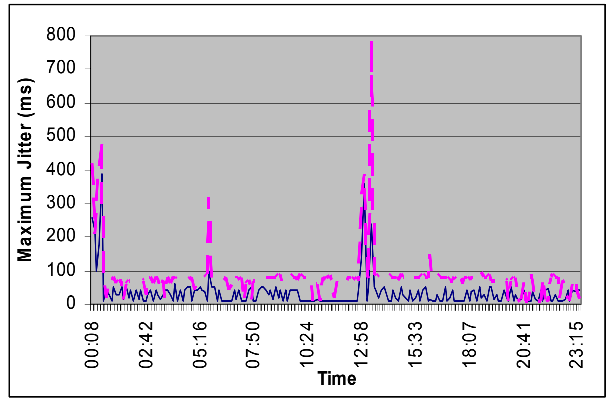

Figure 26 – Maximum Jitter (Single Call)

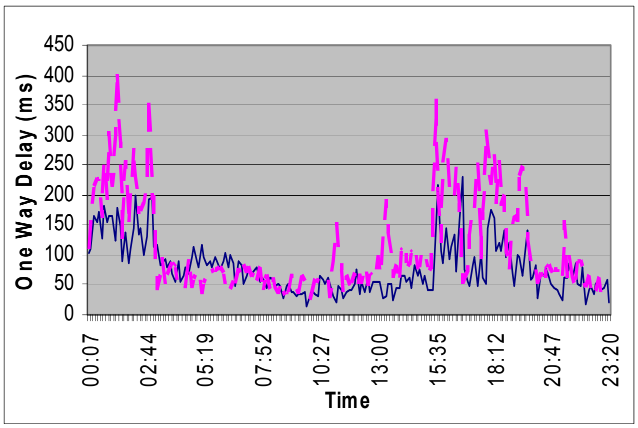

Figure 24 - One Way Delay G711u (1st of 2 Calls)

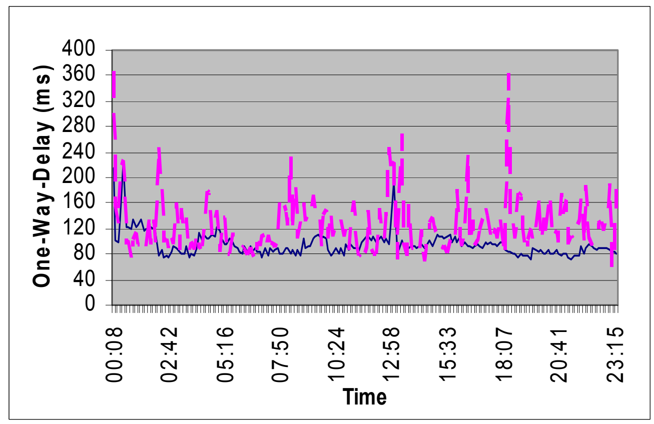

Figure 27 - One Way Delay G711u (Single Call)

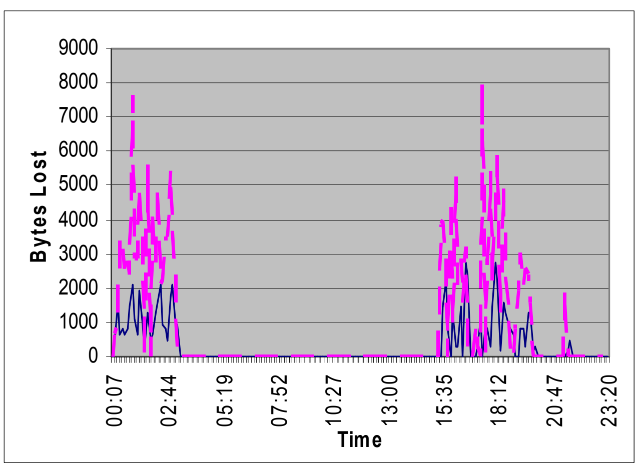

Figure 25 – Bytes Lost G711u (1st of 2 Calls)

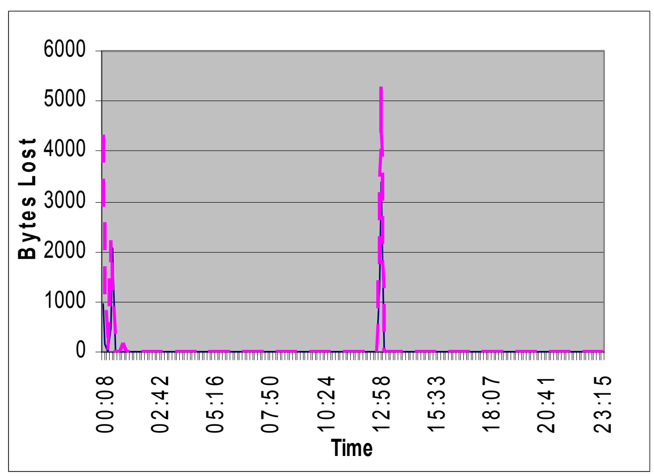

Figure 28 - Bytes Lost G711u (Single Call)

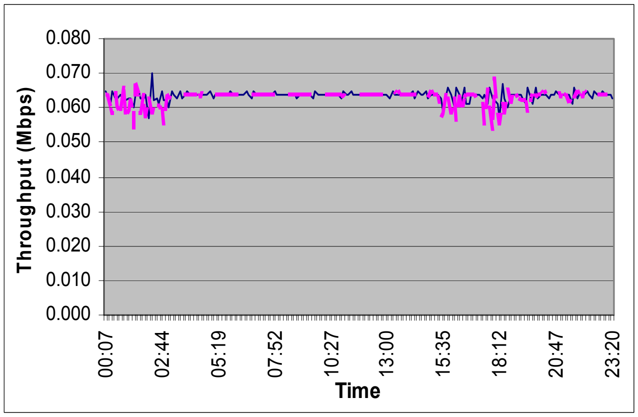

Figure 29 – Throughput G711u (1st of 2 Calls)

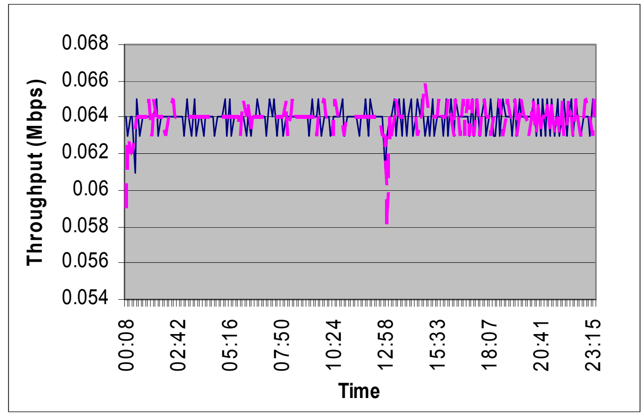

Figure 30 – Throughput G711u (Single Call)

Dow nstream
Upstream

G.729 Codec

The next test emulated a VoIP call using the half duplex G.729 Codec. To emulate the half duplex nature of the codec the test was performed in the upstream and downstream directions independently. For the downstream direction, the call showed an average MOS value of 3.68 as shown in Figure 31. Due to its bit rate, and the way that it converts voice into bits, G.729's theoretical maximum MOS rating is only 4.07. The throughput for the call in this direction held steady at around 8 kbps as shown in Figure 43 and only two brief instances of data loss occurred. The one-way delay in the downstream direction often spiked over 200 ms as shown in Figure 38.

For the upstream direction, the call showed an average MOS value of 2.67 as shown in Figure 34. The main cause of this poor performance was the jitter buffer dropping significant amounts of datagrams as shown in Figure 35. Although, the RFC 1889 jitter shown in Figure 36 never exceeds 18 ms, the non-averaged maximum jitter results in Figure 40 show that some packets throughout the call actually had jitter over 50 ms. Once again the throughput held steady at around 8 kbps and there was only one brief instance of non-jitter buffer related data loss as shown in Figure 42.

As seen in the previous G.711u Codec emulation the downstream portion of the test performed better than the upstream portion of the test. The average MOS values for the G.729 VoIP call are shown in Table 4. The plots of the important statistics recorded during the calls for the downstream and upstream tests are shown side by side below in Figures 31 – 44.

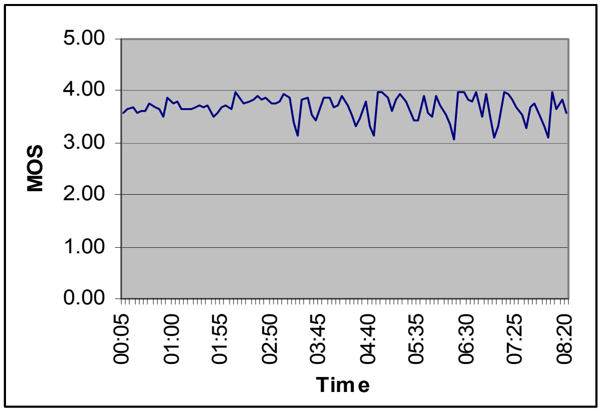

Figure 31 - MOS G729 Downstream

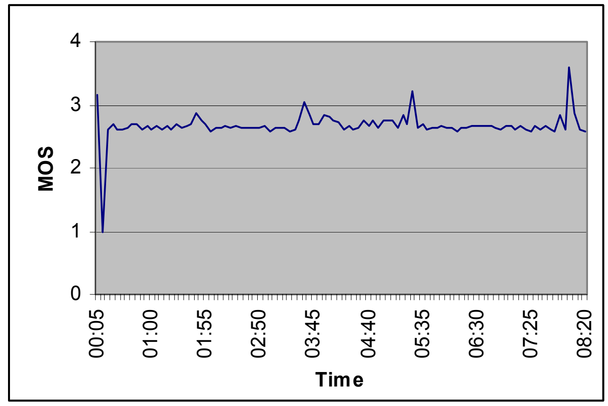

Figure 34 – MOS G729 Upstream

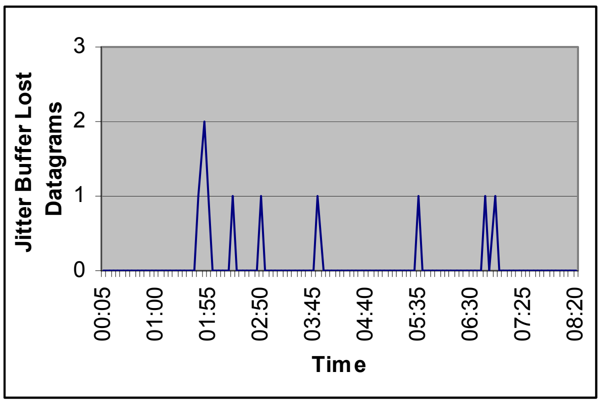

Figure 32 – Jitter Buffer Lost Datagrams (Downstream)

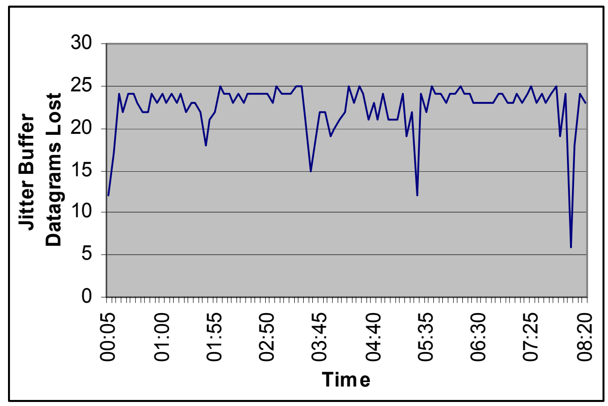

Figure 35 – Jitter Buffer Lost Datagrams (Upstream)

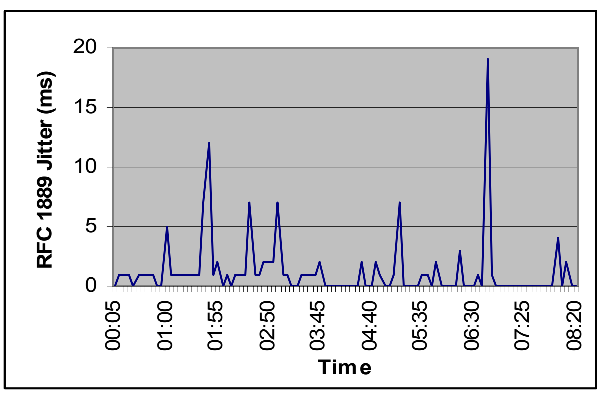

Figure 33 - RFC 1889 Jitter (Downstream)

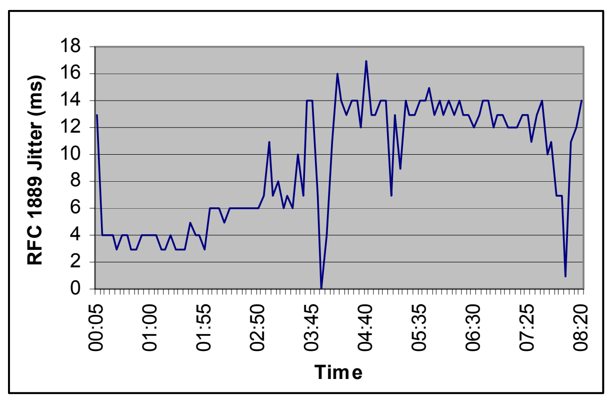

Figure 36 - RFC1889 Jitter (Upstream)

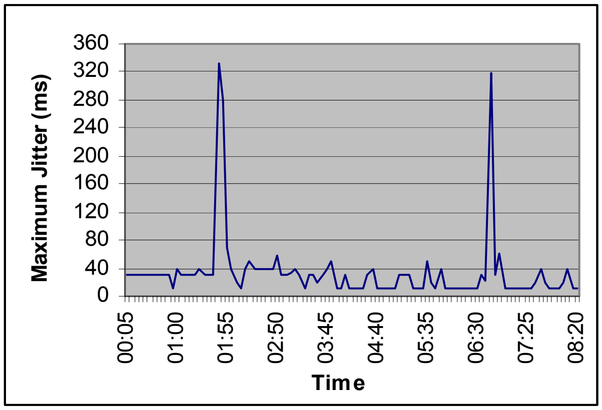

Figure 37 – Maximum Jitter (Downstream)

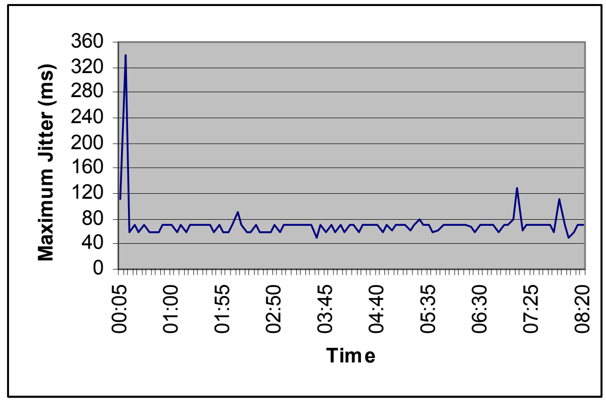

Figure 40 – Maximum Jitter (Upstream)

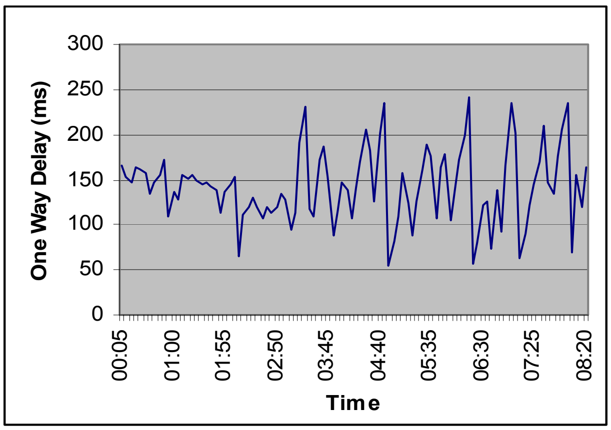

Figure 38 - One Way Delay (Downstream)

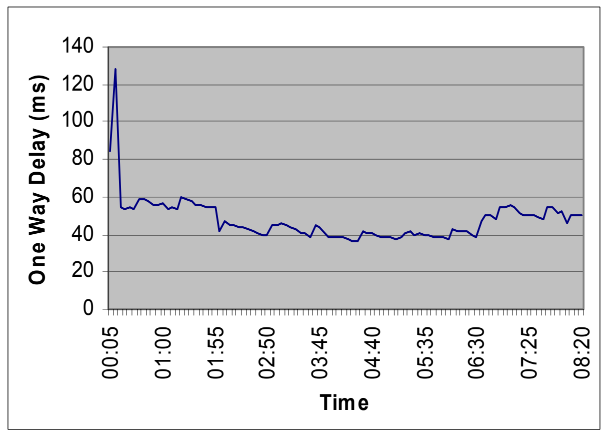

Figure 41 - One Way Delay (Upstream)

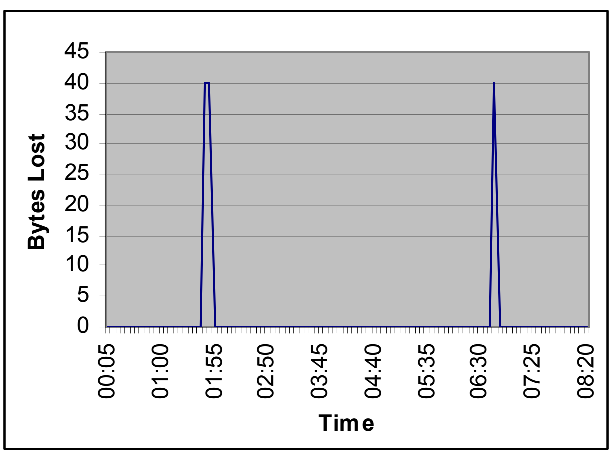

Figure 39 - Bytes Lost (Downstream)

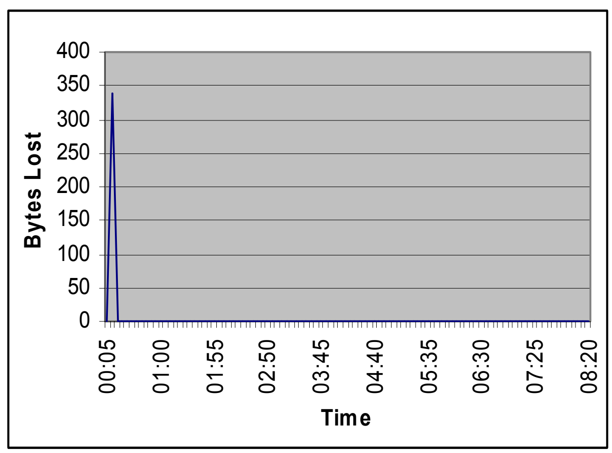

Figure 42 - Bytes Lost (Upstream)

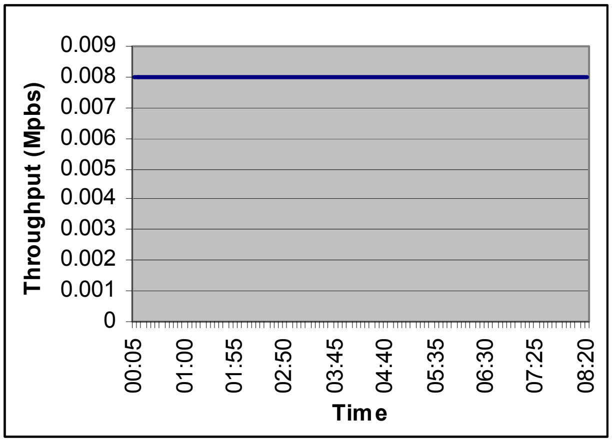

Figure 43 - Throughput (Downstream)

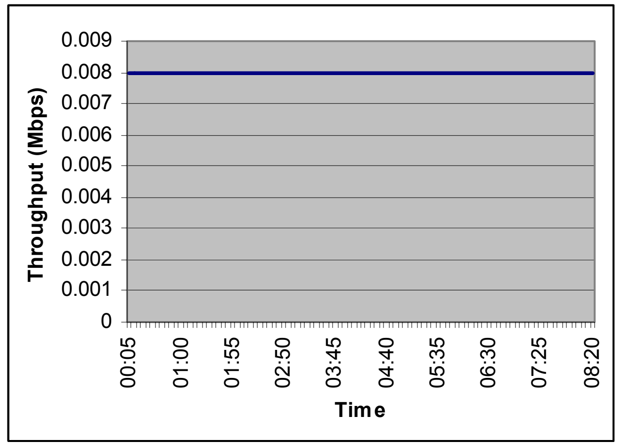

Figure 44 – Throughput (Upstream)

Summary of VoIP Results

The VoIP evaluation performed on AT&Ts WBB link demonstrated that the link as configured could support one full-duplex call using the 711u codec. Additional VoIP calls require multiple modems along with the allocation of single VoIP calls per modem due to low MOS ratings. The lower quality half duplex 729 Codec did not perform well for one call, although it used less bandwidth. This behavior highlighted that although the required throughput was available for each call, other factors such as jitter and one-way delay had a large impact on the quality of the calls. AT&T's WBB technology was not configured or optimized for VoIP and time sensitive traffic. The optimization and configuration could provide enhanced performance to support VoIP traffic by providing better MOS ratings and possibly more calls on the link. AT&T forewarned NCS of this fact before testing, but VoIP was still tested to observe what the call quality would be without optimization.

Table 4 – VoIP MOS Results

Simultaneous Calls	Call #	Direction	Average MOS	
			711u	729
1	1	Downstream	4.18	3.68
		Upstream	3.39	2.67
2	1st Call	Downstream	3.30	N/A
		Upstream	2.40	N/A
	2nd Call	Downstream	3.34	N/A
		Upstream	2.40	N/A

STREAMING MEDIA TESTING

The multimedia streaming test evaluated the quality of an Internet conferencing program, Microsoft's NetMeeting, which provides real-time video and audio conferencing streams. The test emulated a video and audio conference between two people, with traffic flowing simultaneously upstream and downstream. Unlike VoIP, there is no rating system equivalent to MOS that can rate the overall quality of the conference emulated. The video and audio conference stream results were measured against the tolerances shown in Table 5 to determine conference quality. Similar to the VoIP testing, the test first established the quality of a single conference session. After the initial conference, the number of simultaneous conference sessions was increased by one each time until the quality of the conferences no longer met the requirements in Table 3.

Table 5 – Interactive Video Requirements

Metric	Tolerances
One-Way Latency	< = 150 ms
Jitter	< = 30 ms
% of Bytes Lost	< = 1 %

Streaming Media Test Results

One Audio / Video Conference

The results of the emulation of the video and audio conferencing program NetMeeting were very similar to the results of the VoIP tests. This behavior was expected given that VoIP and video and audio conferencing are both real-time applications and are susceptible to the same performance degradation factors. The results demonstrated that the WBB technology in its current configuration could handle a single video and audio conference session and stay within the requirements listed in Table 5. As seen in Figures 45 and 47 a single conference did not experience any significant data loss. As shown in Figures 46, 48, 49, and 51 the downstream and upstream jitter and one way delay measurements remained below the thresholds outlined in table 3. As expected, due to its larger bandwidth, the downstream portions of the conferences performed better than the upstream portions. The downstream traffic experienced less jitter (Figure 46 and Figure 48) and one-way delay (Figure 49 and Figure 51). The throughput for the calls held steady at 64kbps for video (Figure 50) and 12kbps for audio (Figure 52).

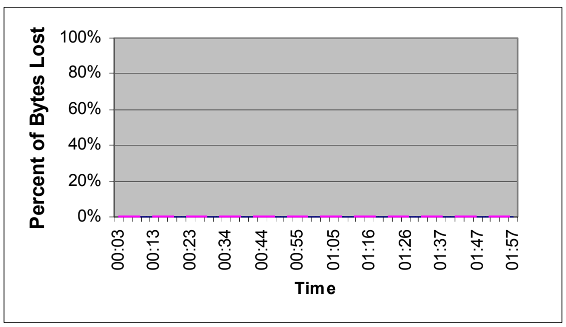

Figure 45 – Percent of Bytes Lost (Video)

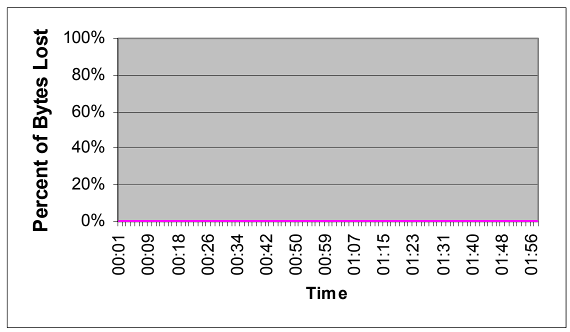

Figure 47 – Percent of Bytes Lost (Audio)

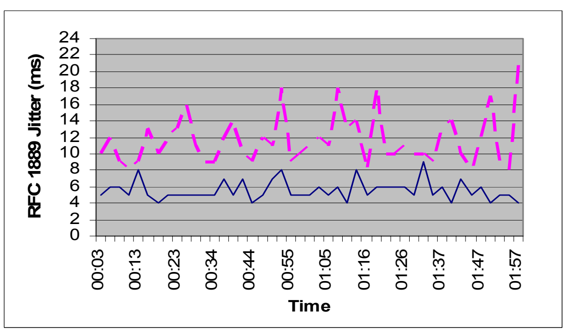

Figure 46 – RFC 1889 Jitter (Video)

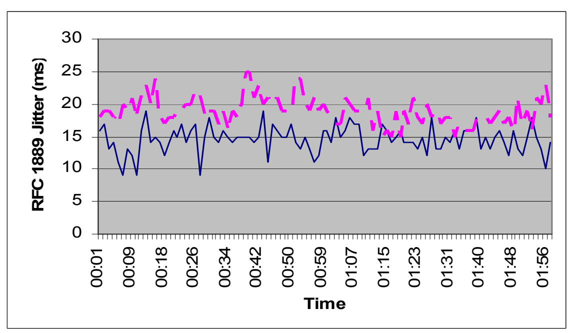

Figure 48 – RFC 1889 Jitter (Audio)

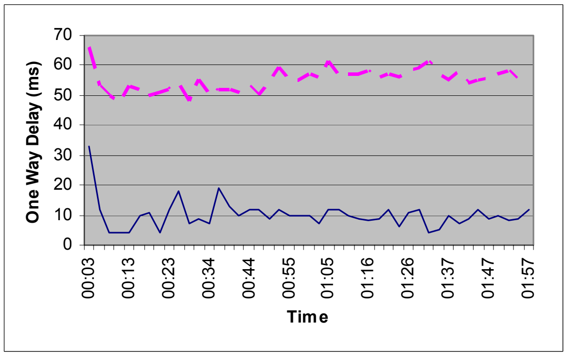

Figure 49 – One Way Delay (Video)

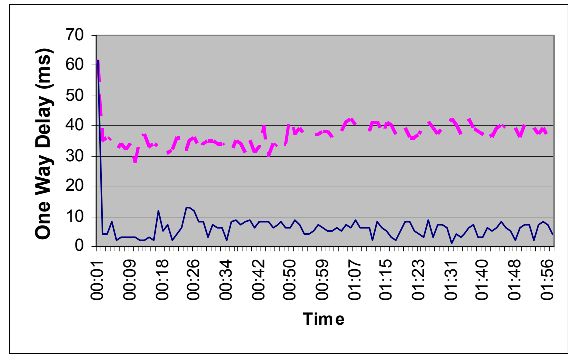

Figure 51 – One Way Delay (Audio)

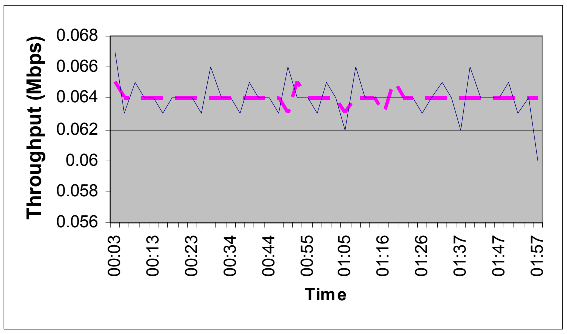

Figure 50 – Throughput (Video)

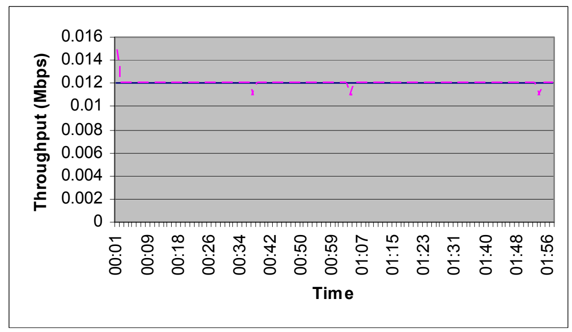

Figure 52 – Throughput (Audio)

Two Audio / Video Conferences

When emulating two simultaneous video / audio conference sessions across the link, the performances of both conference sessions degraded significantly in comparison to the single conference session. The cause of the poor performance can be attributed to data loss. As shown in Figures 52 and 53, the percent data loss measurements were often above the 1% requirement. Jitter (Figure 55 and Figure 56) and one way delay (Figure 57 and Figure 58) also spiked above their acceptable thresholds of 30ms and 150ms respectively. However, during the periods where data loss was not present, the conference sessions did stay within the requirements outlined in Table 5.

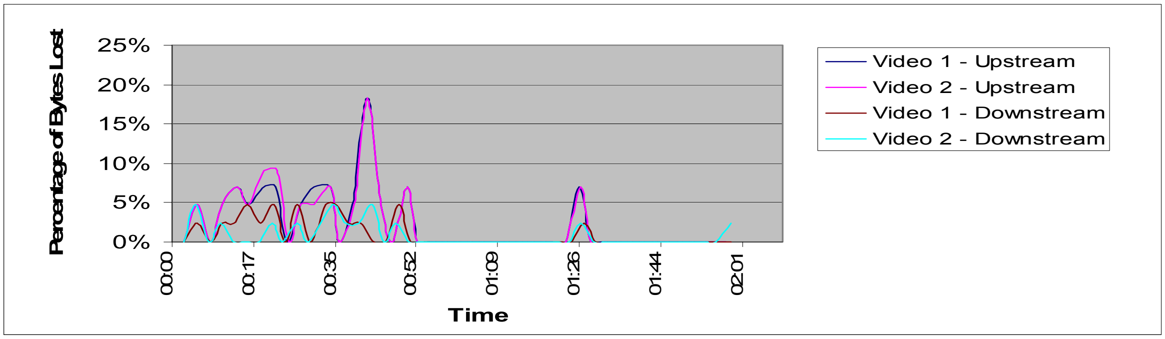

Figure 53 – Percent of Bytes Lost – 2 conferences (Video)

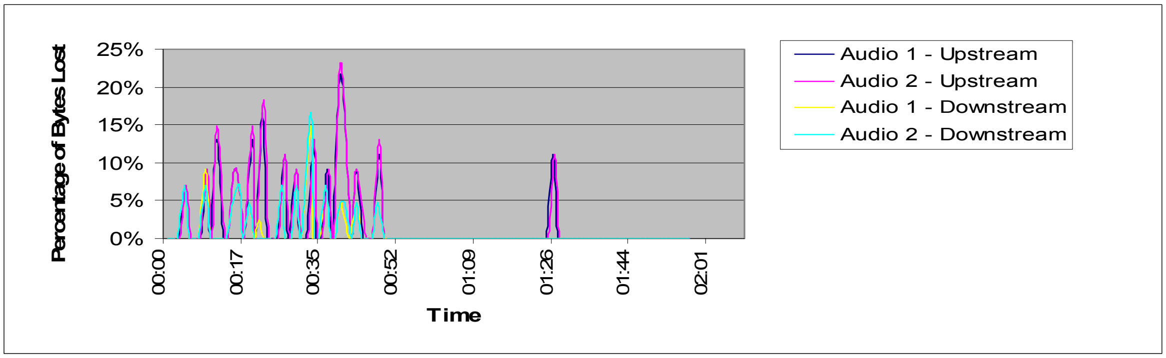

Figure 54 – Percentage of Bytes Lost - 2 Conferences (Audio)

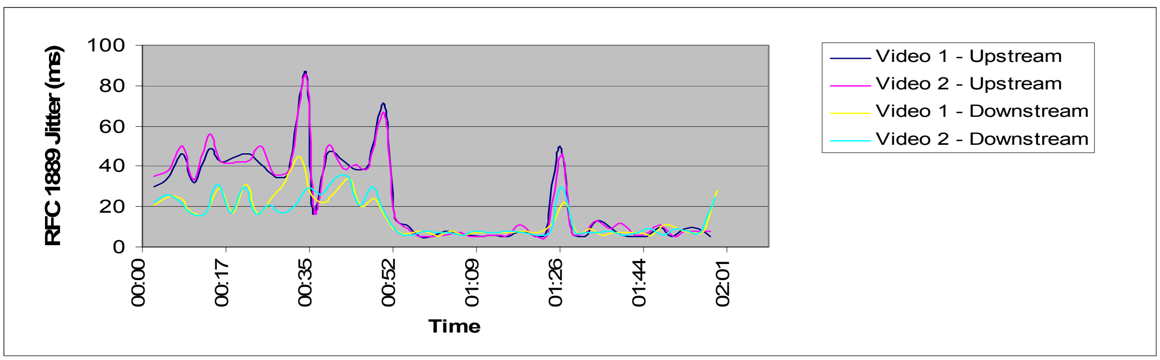

Figure 55 – RFC 1889 Jitter – 2 Conferences (Video)

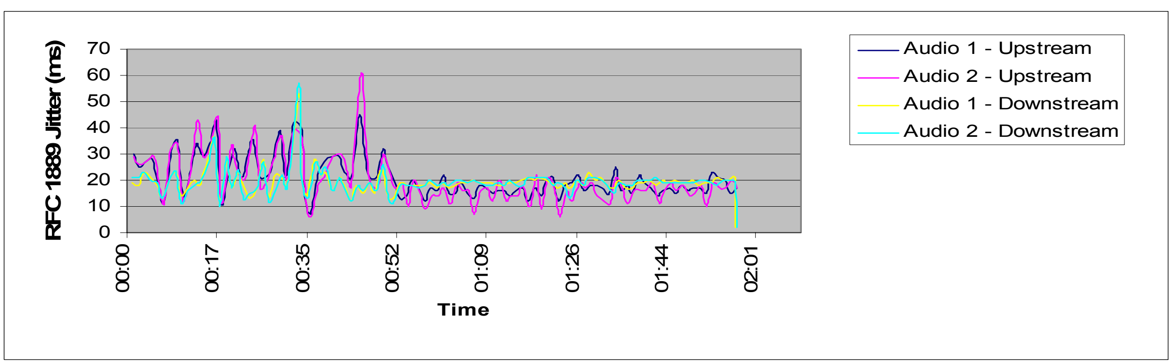

Figure 56 - RFC 1889 Jitter – 2 Conferences (Audio)

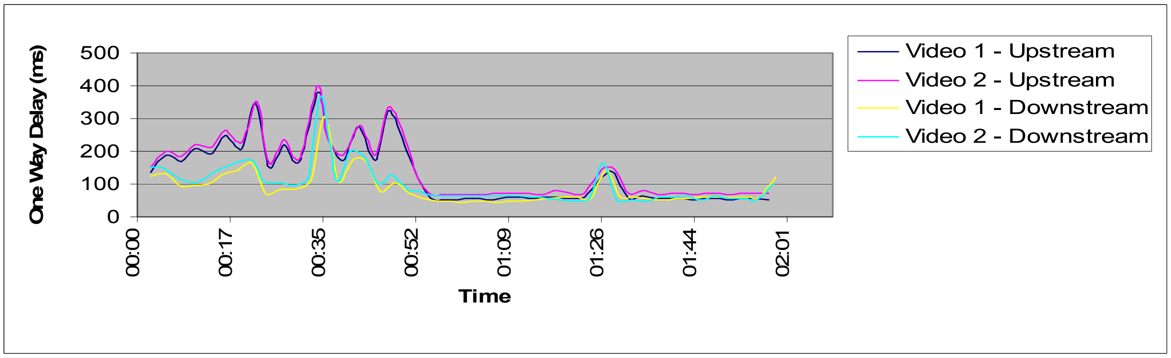

Figure 57 - One Way Delay – 2 Conferences (Video)

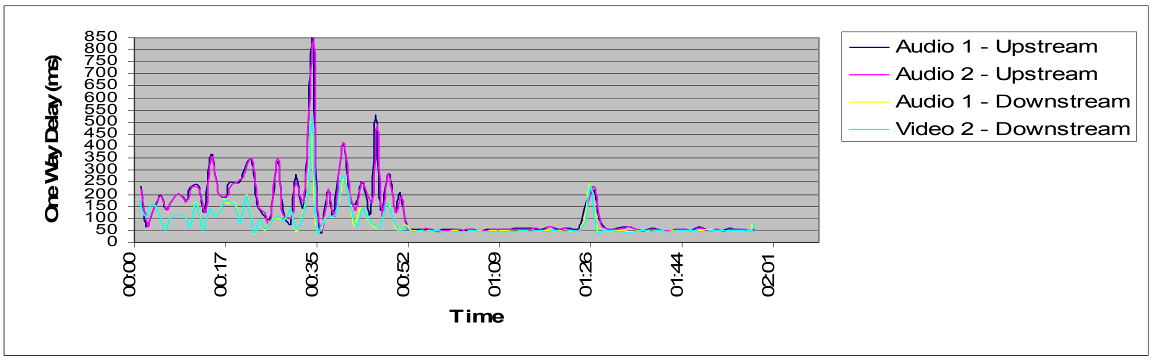

Figure 58 – One Way Delay – 2 Conferences (Audio)

APPENDIX D- TEST EQUIPMENT DETAILS

IXIA hardware and software generate stateful and stateless traffic flows across the link and capture results for performance evaluation. The IXIA Ethernet modules provide wire-speed traffic generation and analysis, upper layer protocol emulation running simultaneously with wire-speed packet generation, and capture various performance data statistics on layer 2 – Data Link and Layer 3 – Network Routing. The IxChariot software is a traffic pattern analysis and decision support tool emulating real world application data. It offers thorough application assessment and device testing by emulating hundreds of protocols across network endpoints and providing upper layer session emulation. The integration of IXIA hardware and IxChariot software provides the ability to simultaneously transmit stateless and stateful traffic sessions and measure performance across all layers of the protocol stack. Additional tools and software are used to monitor and collect data for further analysis and correlation as described below.

IXIA HARDWARE AND SOFTWARE

IXIA's scalable test solutions generate, capture, characterize, and emulate network and application traffic, establishing definitive performance metrics of the applications. IXIA's integrated hardware / software platforms can generate and analyze Layer 2 and 3 network traffic, while simultaneously testing Layer 4 through 7 application traffic. The Wireless Broadband links will be evaluated using the IXIA 400T chassis, Gigabit Ethernet Load Modules, IxExplorer software, and IxChariot application test software.

IXIA 400T Chassis

The IXIA series of chassis deliver comprehensive solutions for high performance, multi-port traffic generation and analysis. The IXIA chassis share a highly scalable architecture that support daisy-chaining of multiple chassis' for timing synchronization. Each chassis supports Packet Over SONET, 10 Gigabit Ethernet, Gigabit Ethernet, 10/100 Ethernet, and ATM interfaces. The integrated PC, running Windows 2000, is used for configuration and statistics gathering and is able to generate and analyze traffic synchronously over multiple ports simultaneously.

Ethernet Load Modules

IXIA's family of Ethernet load modules offers wire-speed Layer 2 and 3 traffic generation and analysis, scalable routing protocol emulation, and Layer 4 through Layer 7 session emulation. Every port on the load module contains a powerful RISC processor running the Linux operating system and a full TCP/IP stack. This architecture gives flexibility when testing systems or devices that require upper layer protocol emulation running simultaneously with wire-speed packet generation. Traffic is generated in real-time by logic implemented on each IXIA port. The packet streams transmit engine allows generation of up to 256 unique streams on each port. Within each stream, millions of packets can be configured with completely customizable characteristics for every packet header field. Customizable payload contents can also be defined. Frame size can be fixed, vary according to a pattern, or be randomly assigned across a weighted range.

These modules provide extensive statistics and logging capabilities depending on the test configuration, including data capture, data integrity, and sequence checking. Each Gigabit port is configured with 8 MB of capture memory, which can store tens of thousands of packets in real-time. The capture buffer can be configured to store packets based on user-defined trigger and filter conditions. To validate performance, the data integrity function allows packet payload contents to be verified with a unique CRC that is independent of the packet CRC. This ensures that the payload is not disturbed as the header fields are changed and transmitted across the network. Sequence numbers can be inserted at a user-defined offset in the payload of each transmitted packet. Upon receipt of the packets through the system or device under test, out-of-sequence errors are reported in real-time at wire-speed rates. Sequence error thresholds can be defined and measured to distinguish between small, big, reversed, and total errors.

IxExplorer

IXIA's IxExplorer, a Microsoft Windows based software interface, provides a powerful and sophisticated Graphical User Interface (GUI) to manage IXIA test hardware. Each test port can be independently configured with parameters to define streams, filters, and capture capabilities. Users can generate customized traffic patterns using packet streams and easily configure packets with incrementing, decrementing, or random MAC and IP addresses and through user-defined fields. In addition, traffic for various network protocols can be customized, transmitted, and received on each port. Comprehensive statistics allow users to perform an in-depth analysis of the performance of the system or device under test using tables and graphical views.

IxChariot

The IxChariot software is a traffic pattern analysis and decision support tool emulating real world application data. Incorporating the IxChariot console and performance endpoints, IxChariot offers thorough application assessment and device testing by emulating hundreds of protocols (TCP, UDP, RTP, SPX, IPX, SNA, IP Multicast, etc.) across thousands of network endpoints. IxChariot provides the ability to predict the expected performance characteristics of any application running on wired and wireless networks. Over a hundred pre-programmed scripts are capable of emulating all common enterprise protocols and services including VoIP, Multicast, and IPv6 applications running on thousands of endpoint pairs. The software provides the ability to create sophisticated traffic pattern, throughput, jitter, lost data, QoS, and latency analysis using performance endpoints running on various operating systems.

IxChariot sends stateful traffic, which maintains the context of the session between the sender and receiver in order to provide some service level guarantee, commonly reliable, in-order delivery. IxExplorer transmits stateless traffic, which are packets without context reservation to any related stream of packets/flows/sessions/protocols or applications.

The IXIA test configuration and test results data will be transmitted across a separate out-of-band connection. This out-of-band network connectivity is over the Internet between the two labs. When a single network is used, both management and test traffic travel over the same network, providing additional load to the network elements being tested. The preferred method uses two networks and serves to isolate test traffic from all other traffic.

CISCO CATALYST SWITCH

The Catalyst 3550 switch is a stackable 10/100 and Gigabit Ethernet switch. The Catalyst switch was used at location Site 1 to build the Local Area Network (LAN) and connect all other devices together via a Fast Ethernet connection. It allows a central point for managing all hardware and gathering test and performance data. The Gigabit Ethernet ports may be used for additional network connectivity for further testing / demonstrations.

LAPTOP

The IBM ThinkPad T40 laptop operates on the Windows 2000 Operating System (OS). The laptop was used to make network configuration changes, manage various network-monitoring applications, and capture test and performance data as stated in this test plan.

VIRTUAL NETWORK COMPUTING (VNC) SOFTWARE

Virtual Network Computing (VNC) is remote control software that allows viewing and interacting with one computer (the "server") using a simple program (the "viewer") on another computer anywhere on the Internet or via remote dial connectivity. VNC enables remote management of the laptops at each location for configuration and management of the IXIA endpoints during the evaluation.

APPENDIX E - ADDITIONAL APPLICATION INFORMATION

VOICE OVER IP

VoIP traffic consists of Real-Time Transport Protocol (RTP) streams in each direction. RTP is designed to send data in one direction with no acknowledgement. The header of each RTP datagram contains payload type (which Codec to use), sequence number (helps receiver reassemble data and detect missing or out-of-order datagrams), time stamp (enables receiver to reconstruct the timing of the original data), and source ID (enables receiver to distinguish multiple, simultaneous streams).

Mean Opinion Score Estimation

Several statistics will be collected for VoIP performance. One characteristic unique to VoIP is the mean opinion score (MOS). In most cases, MOS is based on subjective opinions from testing participants. For testing, MOS will be calculated by IXIA using a modified version of the International Telecommunications Union (ITU) G.107 standard E-model equation, which evaluates the quality of a transmission by factoring in the "mouth-to-ear" characteristics of a speech path. Then it calculates an R-value, which correlates directly with the MOS estimate. IXChariot modifies the E-model slightly and uses the factors shown in Table 6 to calculate the R-value and the MOS estimate.

Table 6 - Estimated Mean Opinion Score Components

One-Way (Network) Delay	Similar to propagation delay, except the delay factors associated with the network itself are included. IxChariot measures this by synchronizing the end-points' timers and determining delay in a single direction.
End-to-End Delay	Latency measured by adding the following factors: one-way delay, packetization delay, jitter buffer delay, and additional fixed delay. As with one-way delay, end-to-end delay is in a single direction between the end-points, but it extends to the VoIP HW to include all delay factors. One-way delay measures only network delay
Packetization Delay	Delay associated with conversion of the voice signal from analog to digital is factored into the score calculation. It varies according to the Codec type being used.
Jitter Buffer Delay	Delay associated with jitter buffers that work to reduce variability in datagram inter arrival times.
Additional Fixed Delay	Delay factor from a known, constant source of delay. Necessary for accuracy if the test emulates HW associated with a fixed delay factor.
Data Loss	Total number of datagrams lost. When a datagram is lost, you can lose an entire syllable, and the more datagrams lost consecutively, the more that clarity suffers. IxChariot factors in lost data and includes the amount of consecutive datagram loss that was measured.
Jitter Buffer Lost Datagrams	Number of datagrams lost due to jitter buffer overruns and under-runs.

A MOS of 5 is excellent; a MOS of 1 is unacceptably bad. Table 7, shown below from ITU G.107, summarizes the relationship between the MOS and user satisfaction.

Table 7 - ITU G.107 MOS Scale

Mean Opinion Score (lower limit)	User Satisfaction
4.34	Very satisfied
4.03	Satisfied
3.60	Some users dissatisfied
3.10	Many users dissatisfied
2.58	Nearly all users dissatisfied

STREAMING MEDIA

Streaming scripts emulate multimedia applications that send data without acknowledgements. Datagrams are sent in only one direction, from E1 to E2. Of the many IXIA streaming scripts, NetMeeting will be used with the Real-Time Transport Protocol (RTP). When running a streaming script, E2 keeps statistics on lost data and returns this information as part of the results. In addition, when using RTP, E2 records statistics on jitter.

Typical multimedia applications use various packet sizes. End-point multimedia support uses a 12-byte header for RTP. In addition to these headers, the protocol stack adds 8 bytes for the User Datagram Protocol (UDP) and 20 bytes for IP.

APPENDIX F - METRICS DEFINITIONS

Listed below are definitions of the metrics recorded for each of the tests.

FTP

Throughput – The average, minimum, and maximum throughput values were recorded for each test.

HTTP

Response time - the time, in seconds, needed for one transaction

EMAIL

Response time - the time, in seconds, needed for one transaction

MULTIPLE SERVICES

Response time - the time, in seconds, needed for one transaction

VOIP

Throughput – The average, minimum, and maximum throughput values were recorded for each test.

Mean Opinion Score (MOS) – An ITU standard derived from actual listeners' subjective judgments about call quality

One Way Delay (Latency) - Similar to propagation delay, except the delay factors associated with the network itself are included. IxChariot measures this by synchronizing the end-points' timers and determining delay in a single direction.

RFC 1889 Jitter – RFC 1889 jitter (calculated for all RTP pairs) shows mean statistical deviance of packet inter-arrival times over a period of time.

Bytes Lost – The amount of Data expressed in bytes sent by Endpoint 1 that never reached Endpoint 2.

Jitter Buffer Lost Datagrams - Datagrams lost due to jitter buffer overruns, or the number of datagrams that had a delay variation greater than the jitter buffer size, and jitter buffer underruns, datagrams that arrived too quickly while the jitter buffer was still full.

Streaming Media

Throughput - Throughput – The average, minimum, and maximum throughput values were recorded for each test.

One Way Delay (Latency) - is the difference between the time a datagram was sent by E1 and the time it was received by E2. Delay is expressed as an average for all datagrams sent in a single timing record.

RFC 1889 Jitter – RFC 1889 jitter (calculated for all RTP pairs) shows mean statistical deviance of packet interarrival times over a period of time.

Percent of Bytes Lost – Data sent by Endpoint 1 that never reached Endpoint 2, expressed as a percentage of the total amount of bytes sent.

APPENDIX G – REFERENCE DOCUMENTS

IxExplorer Users Guide: General Configuration and Operation, Release 3.70; Ixia, Part No. 909-0002-01, Rev. A, August 2003.

IxChariot Users Guide, Release 5.0; Ixia, Part No. 909-0154, Rev. A, January 2004.

IxChariot Application Scripts, Release 5.0; Ixia, Part No. 909-0157, Rev. A, January 2004.

APPENDIX H - ROUTE DIVERSITY PROJECT BACKGROUND

On September 11, 2001, terrorists struck the World Trade Center in New York City and the Pentagon near Washington, DC. Reports indicated that telecommunication assets near the affected areas were either congested or incapacitated, causing users to experience intermittent or no voice service. The reports of these events generated concern among White House officials that key federal agencies in Washington, DC might lose critical wire line voice and data communications services if the infrastructure was damaged or destroyed.

To highlight the concerns of the White House officials, the National Security Council (NSC) raised the issue of telecommunications resiliency in its meeting of October 5, 2001, stating—

> *"Key federal agencies may be at risk of losing wire line communications services in certain emergencies where telecommunications infrastructure gets damaged or destroyed."*[3]

The NCS, which is responsible for ensuring National Security and Emergency Preparedness (NS/EP) communications in times of network congestion or outage, addressed this concern by investigating the possibility of a route diversity capability for federal agencies. The NCS established the RDP (formerly known as the Backup Dial Tone Project) and took the following steps:

Evaluated the need for a route diversity capability in the Washington, DC area and determined whether such a capability would have been helpful in the New York City and Washington, DC areas on September 11, 2001 (Phase I—RDP).

Evaluated various technical approaches to providing such a capability (Phase II—RDP)

Determined the cost and schedule for deploying chosen technical approaches (Phase III—RDP; ongoing).

During these initial phases, the RDP has proven to be a constructive, valuable program in helping to address the NSC directive. The CRM, which is Phase IV of the RDP, uses lessons learned from the first three phases to provide individual federal agencies with a tool to evaluate and address RDP capabilities.

Methodology of RDP Project

The NCS approached the RDP initiative in the following phases:

- **Phase I: Identify Key Federal Facilities and Perform Generic Analyses—**compiled a comprehensive list of potential vulnerabilities based on generic government facility architectures and a list of technical solutions for addressing these vulnerabilities. This phase was completed in January 2002.

[3] National Security Council Memorandum. Subject: Minutes from October 5, 2001, Meeting on Selected NS/EP Telecommunications Projects (U)

- **Phase II: Narrow the Solutions and Interview Federal Agency Representatives—** refined this list of solutions, relying on interviews with government agencies to enhance the generic architectures with real-world data. This phase was completed in August 2002.
- **Phase III: Conduct Technology Demonstrations—**selected the final technologies for demonstration, determined potential agency test sites, and developed the test demonstrations for evaluation. This phase is ongoing, and includes the demonstration detailed in this report.
- **Phase IV: Develop Route Diversity Methodology—**applies knowledge gained from the first three phases to develop a methodology for assessing route diversity at a single agency location. The methodology highlights areas that may be weak in route diversity in an agency's communications infrastructure. After the weak areas are highlighted, potential mitigation solutions that were identified and evaluated in previous phases can be reviewed for implementation at a specific site. This phase is ongoing.
- **Phase V: Network Resiliency and Emergency Support Functions #2 (ESF #2) Support** - includes developing tools and leveraging previous RDP efforts to assist member agencies in meeting legislative and regulatory requirements pertaining to route diversity and ESF #2 needs. This phase is ongoing.