H.323 Beacon: An H.323 Application Related End-to-End Performance Troubleshooting Tool *

Prasad Calyam,Weiping Mandrawa, Mukundan Sridharan, Arif Khan, Paul Schopis
OARnet, 1224 Kinnear Road, Columbus, Ohio 43212
{pcalyam,wmandraw,sridhara,akhan,pschopis}@oar.net

ABSTRACT

H.323 protocol based Voice and Video conferencing solutions are established popular technologies both in industry and academia. However, these bandwidth intensive applications are often plagued by various performance problems. In this paper we describe a few common end-to-end performance problems noticed over several years in our ITU-T H.323 protocol-based "Video and Voice over IP" (VVoIP) infrastructure and provide a tool for troubleshooting these performance issues. Based on our operations experiences, we are developing a tool called the "H.323 Beacon" to measure, monitor and qualify the performance of H.323 sessions. We describe the development architecture and feature set of the H.323 Beacon, which can be used by a novice end-user, network engineer or conference-operator using VVoIP systems. H.323 Beacon is an application-specific network measurement tool that provides H.323-protocol specific evidence and other information necessary to troubleshoot H.323 application performance in the network and at the host (end-to-end). We also present 2 use-cases for the H.323 Beacon that demonstrate the utility of the tool for VVoIP performance debugging. Finally, we outline essential best practices for prevention and efficient resolution of intermittent and imminent failures of H.323 VVoIP applications.

Categories and Subject Descriptors

H.4.3 [**Information Systems Applications**]: Communications Applications; C.4 [**Computer Systems Organization**]: Performance of Systems

General Terms

Measurement, Performance, Human factors

Keywords

H.323, Videoconferencing, VoIP, Network Measurement

*This work was supported in part by The Ohio Board of Regents and Internet2

SIGCOMM'04 Workshops, Aug. 30+Sept. 3, 2004, Portland, Oregon, USA.

1. INTRODUCTION

H.323 [7] is an umbrella standard from the International Telecommunications Union (ITU) for developing real-time multimedia communication applications, such as audio and video conferencing, over packet-switched networks (Internet). Internet-based audio and video conferencing has established itself as an important technology in supporting collaboration both in academia and industry.

Having provided VVoIP services for many years and being a Global H.323 Network Operations Center (NOC) for the Internet2 Commons and the Megaconferences [5], we have closely observed rapid advancements in H.323 related technologies in areas such as interoperability, supplementary services, protocol extensions and cost-effectiveness. We have also encountered and resolved numerous end-to-end performance problems affecting H.323 protocol based VVoIP applications.

Since H.323 is a recent technology, there is a significant dearth of affordable and relevant diagnostic tools available to end-users, network engineers or conference operators to identify and troubleshoot end-to-end performance problems related to H.323 applications [6, 11, 1, 4]. Also, because the H.323 protocol has many inherent idiosyncrasies, many of the popular network troubleshooting tools, such as ping and traceroute, do not suffice in the context of H.323 application performance troubleshooting in the Internet.

Based on our experience with many of the aforementioned tools, we have developed the "H.323 Beacon" which has all the essential tool features available in some popular tools and also has additional features that we have developed, that we have found to be vital for H.323 application performance troubleshooting. To demonstrate the utility of the H.323 Beacon to effectively identify and resolve end-to-end performance issues affecting H.323 applications, we discuss two use cases that feature commonly occurring problem scenarios we have observed at our NOC. The use cases contain our analysis process and conclusions based on the various data points, such as remote site network setup and network path measurement data, collected during the troubleshooting tests. Finally, we outline end-user recommendations we have successfully adopted at our NOC to minimize the expenditure of valuable network engineering resources, to prevent the occurrence of the most common problems and to minimize time for problem resolution.

The reminder of the paper is organized as follows: Section 2 describes common end-to-end performance problems affecting VVoIP applications, Section 3 discusses the H.323 Beacon tool architecture and various feature sets, Section 4 presents two use cases that demonstrate the utility of the H.323 Beacon tool in VVoIP performance debugging, Section 5 outlines essential best practices for prevention and efficient resolution of intermittent and imminent failures of H.323 VVoIP applications and Section 6 concludes the paper.

2. COMMON H.323 APPLICATION PERFORMANCE PROBLEMS

When faced with a performance issues concerning an H.323 application, in routine operations within our NOC, our problem resolution frequently involves identifying a last-mile bottleneck. Most Network backbone infrastructures, especially in academic communities, have been observed to be stable and highly reliable. Problems arising at last-mile links can be attributed to the frequent day-to-day variations in end-user and network resource management patterns and in a few instances, due to device misconfiguration by inexperienced networking personnel.

At the core, a VVoIP system constitutes three types of devices: end-point devices, network devices and application-service devices. Network devices include routers and switches. End-point devices include videoconferencing clients with audio and video handling hardware/software solutions. Application service devices include firewalls, Network Address Translators (NATs), Multipoint Control Units (MCUs), gatekeepers, gateways and other devices that provide end-point and network device interfaces with various services such as security policies, multiple video/audio streams multiplexing and resource management. In the following subsections we divide the common end-to-end H.323 application performance problems under each VVoIP system device category.

2.1 Common problems involving end-point devices

The common problems in this category are: failure of Internet connectivity, failure of audio and video hardware, faulty connections and configurations of audio and video interfaces, lack of lip-syncronisation, misconfigured jitter buffer sizes, reception of low frame rates from remote site, outdated or buggy end-point application software, non interoperable end-point application software, using inferior and inexpensive codecs vs expensive and advanced codecs, lack of forward error correction mechanisms, lack of echo cancellation mechanisms, lack of end-user training and faulty cables.

2.2 Common problems involving network devices

The common problems in this category are: insufficient network capacity or bandwidth in the LAN and WAN for handling multiple high data rate videoconferencing calls, excessive delay, loss, jitter, out of order packets and re-ordered packets in the network, duplex-match, traffic congestion at peak-usage periods of the network, misconfigured or poor priorities for real-time audio and video traffic streams in the network, asymmetric routing with excessive delays on one path, lack of network engineering resource personnel and faulty cables.

2.3 Common problems involving application-service devices

The common problems in this category are: misconfigured firewalls that block required ports, non-H.323 friendly NATs, misconfigured MCUs, misconfigured gatekeeper and misconfigured gateway devices, outdated or buggy application-service device application software.

It is obvious from the above that problems at any device section in the overall end-to-end path coupled with the complexity and scale of the VVoIP system can significantly cause performance degradation that affects end-user experience of the H.323 application. However, since impact levels vary, diagnostic tools aim at identification and resolution of the obvious and high-impact problems. There are numerous technologies and recommendations from standard bodies such as IETF and ITU-T that aim at addressing various performance measurement issues relating to end-user experiences and network metrics that characterize network health in the VVoIP system operations. Middle-box detection techniques to identify essential port blocks, NAT presence detection techniques and other MCU-based multiple session monitoring techniques and other technologies have become commonplace to identify and resolve many of the performance problems previously stated. These technologies have been incorporated by many vendors who provide comprehensive, complex and expensive troubleshooting solutions. However, considering the fact that most problems occur in last-mile sites that often lack the monetary resources and trained staff necessary to invest in expensive troubleshooting tools, an easy to install and use open-source solution such as the H.323 Beacon becomes pertinent and necessary.

3. H.323 BEACON TOOL

3.1 Background

The H.323 Beacon is one of the new-generation network measurement tools that emulate the application itself to troubleshoot end-to-end application performance problems. It can be envisioned as a light weight and handy tool that quickly gives problem indications and information, when an end-user complains of bad experiences with their video or audio conferences. Since the end-user experience of the problem is intermittent and the end-user is typically not as knowledgeable as a network engineer, there is frequent difficulty when communicating their problems to a network engineer. The H.323 Beacon tool tries to address these concerns and addresses isolation and resolution of some of the commonly occurring, high-impact end-to-end problems listed in Section 2.

3.2 Architecture

The H.323 Beacon uses a distributed client/server architecture. The client actually refers to an end-node and the server can be visualized as a core-node. Testing between end-nodes can be achieved by using a number of core-nodes along a test path. This architecture allows the H.323 Beacon to perform end-to-end measurements related to H.323 videoconferencing sessions. Figure 1 shows how the H.323 Beacon can be a part of an advanced network measurement infrastructure that supports various application specific measurements. As seen in Figure 1 the H.323 Beacon Client is a GUI Dialog application that enables users to initiate and analyze

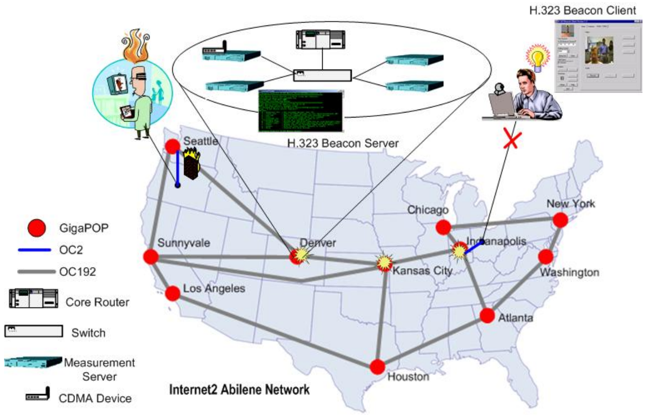

Figure 1: Typical deployment of the H.323 Beacon in an advanced network measurement infrastructure

H.323 test sessions. The H.323 Beacon Server is developed as a console-based application with relevant log generation features for administration purposes, since the typical use of the H.323 Beacon Server does not require excessive human interaction and since general installations are expected to be on servers in measurement racks.

Our goals for developing the H.323 Beacon include developing testing capabilities that can support work flows involving client-to-client, client-to-server and server-to-server tests. An end-user testing to a remote site could use the client-to-server model. A test involving regularly scheduled measurement tests between two servers at strategic points in the network could employ the server-to-server model and when two end-users want to test among themselves they could employ the client-to-client model. In typical videoconferencing operations environments where multiple end-points collaborate via MCUs, a dedicated H.323 Beacon server co-located in the same LAN as the MCUs enables a path test from any end-point to the MCUs.

3.3 Tool Features

The H.323 Beacon uses the open-source OpenH323, PWlib and PTLib libraries [3], on the Windows platform. The core features are based on libraries and sample applications that use the H.323 stack implementation developed by the OpenH323 community. This section lists all of the H.323 Beacon features, many of which have been developed and tested and some of which are currently in development. Please refer to [2] for additional details regarding status of development of the various features. The following paragraphs describe the various feature sets of the tool

3.3.1 Initial Call Setup Failures and Haphazard Disconnections

The H.323 Beacon client has three kinds of test status: "In Session", "Normal Close" and "Exception Close". The "In Session" and "Normal Close" status depict the user initiating the test session and the user ending the test session. The "Exception Close" status depicts a potential performance problem either at the beginning or during the test session that caused the test session to close. The "Exception Close" status is influenced by the code timeouts when there are occurences of signaling failures or media reception failures, which is handled by the OpenH323 libraries to generate various code exceptions. The H.323 Beacon client application intercepts these code exceptions, analyzes them based on the state of the program execution (Ex., failure to initiate signaling contact with remote server, failure of transferring media after successful call setup, etc). It then generates an Alarm that indicates the most probable cause of the problem. A few examples of the alarms are: "No Internet connectivity", "Possible Firewall/NAT presence obstruction", "Excessive network congestion", "Insufficient bandwidth" and "Signaling incompatibility".

3.3.2 Network-path Health Diagnosis

The H.323 Beacon client collects network statistics such as delay, jitter and packet loss in real-time during the test session for the particular test path. The algorithms to determine delay, jitter and packet loss are implemented in the OpenH323 libraries as indicated in the RFC 1889 [10]. At the end of the test session, an Excel file with all the packet timestamps, network performance variations and other relevant information is generated by the H.323 Beacon client. The data in the Excel file can be used to manually generate various kinds of graphs or could be imported into statistical analysis packages. Graphs for network characteristics are generated with colored watermarks for "Good" (green color), "Acceptable" (amber color) and "Poor" (red color) values of delay, jitter and packet loss. The levels for the performance watermarks are set by default to the values obtained from our recent research [8] to enable a typical Beacon-user to visually identify if performance thresholds are being crossed. There are also options to reconfigure these performance thresholds if the Beacon-user chooses to do so.

3.3.3 Audio and Video Reception Quality Assessments for both Local and Remote Ends

When a remote-user complains about the audio or video reception quality of a user, there is a natural desire on the users mind to physically experience the remote users perspective. A novel idea implemented in the H.323 Beacon in the form of the audio/video loop back feature facilitates this. This feature allows recording of the audio and video streams generated by the H.323 Beacon client or even a commercial client, such as the Polycom's Via Video, by the H.323 Beacon server as and when they are received at the remote end. The recorded audio/video files are then sent back to the H323 Beacon Client for play back. The file transfer is done transparently from the end-users perspective ; i.e. a user only needs to push a button to initiate file transfer and observe instant playback! Recording format for audio is the .wav format and the supported formats for video include Avi, Mpeg and QuickTime.

In our current development cycle, we are integrating into the server-to-server module, the E-Model implementation by Leandro et al. [9]. The E-Model is a well established computational model that has been standardized by ITU-T, that uses transmission parameters to predict user perceived voice quality and uses a psycho-acoustic R-scale whose values range from 0 to 100. As shown in Figure 2 the R-values can be mapped to user satisfaction, in terms of Mean Opinion Score (MOS) rankings on a scale of 1-5; 1 being poor and 5 being good. The Emodel has the following components:

$$R = R_0 - I_s - I_d - I_e + A \approx 94 - I_d - I_e$$

Figure 2: Voice Quality Classes

where, R_0 represents the effects of noise and loudness ratio; I_s (Simultaneous Impairment Factor) represents the impairments occurring simultaneously with the speech signal; I_d (Delay Impairment Factor) represents the impairments that are delayed with respect to the speech signal, I_e (Equipment Impairment Factor) represents the impairments caused by transmission equipment and A (Advantage Factor) is used to compensate for the allowance users make for poor quality in return for the ease of access, e.g. cell phone or satellite phone.

The E-Model integration enables deployment of the H.323 Beacon in measurement infrastructures for regularly scheduled tests to determine voice grade capabilities on Internet backbones. Early test results of this module have been used in one of the use-cases described in Section 4.

3.3.4 User-Friendly Features

At the end of the test session, a detailed test session report is generated with information regarding exceptions, summary statistics of network performance measurement data and forward-reverse traceroutes that could potentially be handed over to more experienced support staff to resolve any performance problems experienced by an end-user. The H.323 Beacon client GUI also features user-interface related features such as an activities-log with real-time display of all the activities occurring during the test session and a "Traffic light" functionality to indicate success or failure of the test session.

3.3.5 Options for Customization

To perform customized tests for emulating various scenarios to test network-path health and to manage test results, various Settings options have been provided to the user. These options include: customizing the test session data folder location, TCP/UDP/RTP port settings, H.225 and H.245 configurations, audio/video codec selections for tests, options to save/delete test data and saving the results in various graph formats (bmp/png).

Thus with the above set of features the H.323 Beacon has the potential to troubleshoot various performance problems affecting H.323 applications before, during or even after they arise.

4. USE CASES

In this section we describe 2 use cases which reflect commonly occurring problem scenarios reported at our NOC and our analysis process that resulted in problem isolation and resolution using the H.323 Beacon. In view of maintaining the anonymity of the problem reporting site(s) and to convey our process of troubleshooting in a succinct manner, we use the terminology of NOC End-point and User End-point to refer to the testing sites in the use cases.

4.1 End-user Behind a Firewall

Figure 3: Effect of a firewall on a video stream

It is common to have videoconferencing equipment behind firewalls for security reasons. However, this situation often results in performance degradation that is experienced by the end-user in the form of video and audio impairments, such as video frames freezing and audio dropouts. This occurrence can be attributed to the reduced firewall performance when handling heavy traffic load or can be caused due to misconfiguration of policies.

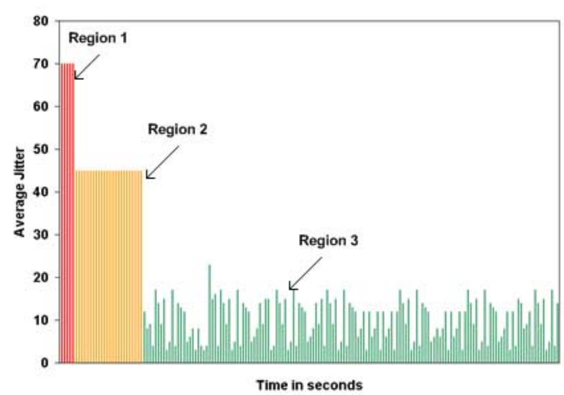

Figure 4: Average Jitter Variations in a video stream behind a misconfigured firewall

Our NOC received a trouble ticket involving a User End-point. The problem description indicated that a few remote sites complained of intermittent video freezing at the User End-point while videoconferencing. No noticeable audio problems were reported. Figure 3 shows a bi-split window showing the User End-point behind a misconfigured firewall. MCU monitoring software indicated an average packet loss of 0.6% and the Tx frame rate of the video to be ∽15 frames per second for a 384Kbps call. Ideally high-quality video consists of ∽30 frames per second. Suspecting a bandwidth limitation for the path, a 128Kbps call was setup and no noticeable loss and reduced Tx frame rate was observed indicating the bandwidth limitation on the remote last mile network. The plot shown in Figure 4 was obtained by running a H.323 Beacon client-to-server test session and plotting the average jitter variations for the test session. The plot shows three distinguishable regions. The first and second regions are the result of a sluggish call setup that affected the entire traffic stream for the duration of the call. The sluggish call set up was caused by the increased delay introduced due to the firewall withholding packets as it verified the misconfigured policies for the new traffic stream. The third region indicates the video streams being finally released for passage with significantly large jitter values to and from the LAN. Also during the test session, H.323 Beacon experienced delayed packet events indicating that the traffic streams were passing through a bottleneck device in the path. Once the bandwidth limitation was mitigated by provisioning additional bandwidth and the firewall misconfiguration problem was resolved, the User End-point was able to overcome their performance problems with their videoconferences.

4.2 An Asymmetric Route Connection

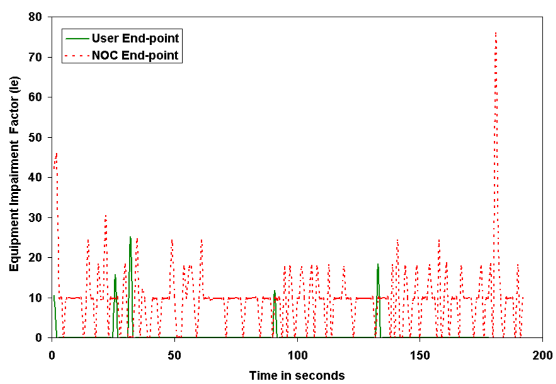

Figure 5: Timeplot of Equipment Impairment Factor

Asymmetric TCP/IP routing has been known to frequently cause performance problems for H.323 applications. The atypical network flows created by asymmetric routes occur when a different interface is used for sending traffic than one that is used to receive traffic. The flows are considered to be unusual because traffic from one end of the connection ($A \longrightarrow B$) travels over a different set of links than does traffic moving in the opposite direction ($B \longrightarrow A$) causing significant delay differences on the paths.

A troubleticket was received at our NOC describing sudden disconnections from videoconferences and also poor audio and video reception at the site of a User. H.323 Beacon server-to-server tests resulted in the plots shown in Figures 5 - 8. Figure 8 shows the MOS variations. The overall MOS obtained at the NOC End-point after analyzing incoming traffic from User End-point was 3.78 and the value at the user End-point side was 4.37. Performing a H.323

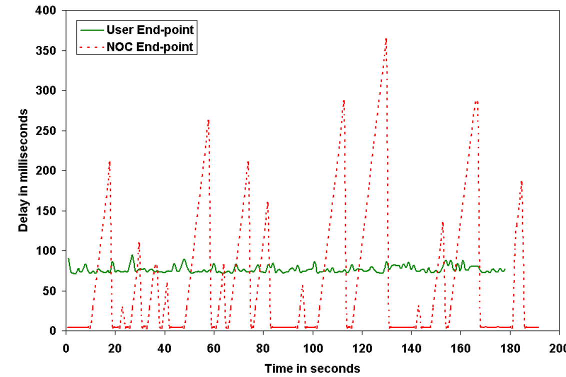

Figure 6: Timeplot of Delay

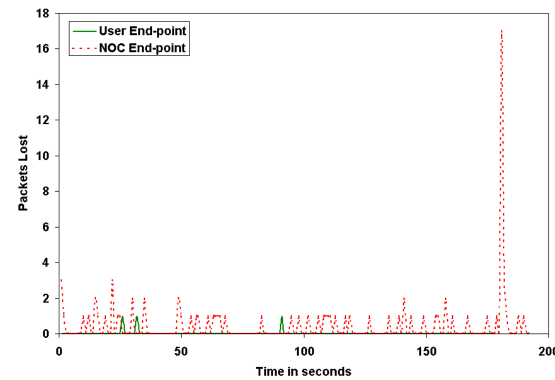

Figure 7: Timeplot of Loss

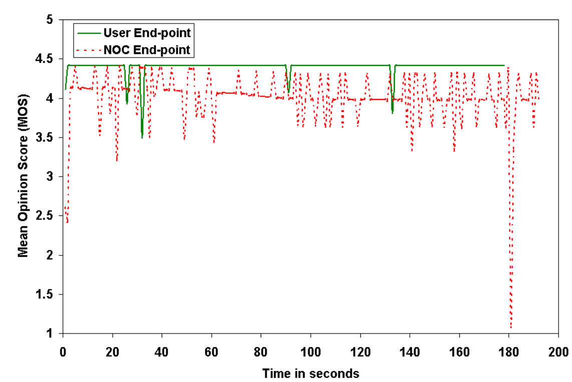

Figure 8: Timeplot of Mean Opinion Score Ranking

Beacon client-to-server test showed significant packet loss values that crossed the "Poor" thresholds in the test session summary graphs. The MCU monitoring software also indicated significant packet loss and out of order packets during videoconferencing sessions with the User End-point. Observing the test session summary report generated by the H.323 Beacon that contained traceroute and reverse traceroute information, an asymmetric routing condition was identified. Based on the routes evidence obtained from the H.323 Beacon tests, the bottleneck link path was identified and a support group at an OARnet's neighbouring ISP was contacted. They acknowledged one of their routers was experiencing a duplex mismatch problem and that they were working toward correcting the problem. On resolution of this issue, further H.323 Beacon tests indicated no apparent performance problems and thus enabled the User End-point to successfully collaborate using videoconferencing in their network.

As shown in the use cases above, creative use of the H.323 Beacon features can help end-users, network engineers and conference operators to identify and resolve many of the common H.323 application performance problems effectively. The H.323 Beacon can also facilitate network measurement research in various H.323 traffic characterization studies as shown in [8]. In [8] the H.323 Beacon was used to predict end-user experience of an H.323 application given any network health diagnostic. Also, it was demonstrated in [8] that a light weight test with the H.323 Beacon actually is reasonably equivalent to a full audio/video test to determine end-user perception of audiovisual quality.

5. BEST PRACTICES

Based on our experience we outline critical recommendations we provide to end-users sites in this section. The following recommendations have been proven to minimize the occurrence of performance problems and foster efficient management of H.323 VVoIP operations.

- Ensure all internal network connections and paths are 100MB and full duplex hard coded from both ends of all routers and/or switches and videoconferencing equipment.

- Prioritize videoconferencing traffic in the network and provision adequate bandwidth expecting that there will to be at least 20% of additional bandwidth overhead during a videoconference. For example, a 384 Kbps call will require atleast 460 Kbps of dedicated bandwidth.

- Ensure that only factory made Cat6 or better cables are used for connections. It is ideal to have cable runs as short as possible and away from electrical fields.

- Populate the LANs used for videoconferencing with programmable switches and routers with adequate processing power; avoid use of Hubs and inexpensive routers that are available in the market.

- Placement of videoconferencing equipment in a Demilitarized Zone (DMZ) and outside of a firewall is highly recommended if possible.

- Have trained technical support staff who have attended trainings such as Site Co-ordinator Training and who understand the basics of preventing and troubleshooting performance problems using relevant tools.

- Keep the software of all videoconferencing and networking equipment updated with the latest software and patches to ensure compatibility and stability.

6. CONCLUSIONS

In this paper, we presented an overview of the common end-to-end performance issues observed in VVoIP environments based on our operational experience. We identified significant techniques and features and developed the H.323 Beacon that could provide an end-user with relevent network and H.323 protocol specific evidence to effectively troubleshoot performance problems. We also discussed 2 use cases that reflect typical problem scenarios and show how the H.323 Beacon could be used to effectively identify and resolve the end-to-end performance problems. Finally we outlined a set of best practices that can be followed at end-user sites to prevent occurrence of performance problems and also to manage VVoIP operations in an efficient and cost effective manner.

The functionalities of the H.323 Beacon tool are being extended on an on-going basis. It is our belief that the H.323 Beacon will supplement the existing open-source and commercial network measurement tools available today in ensuring that trouble free high-quality videoconferencing becomes a reality.

7. ACKNOWLEDGEMENTS

We would like to thank Veena Raghavan, Prashant Reddy, Leandro Lustosa and Paulo Rodrigues for contributing towards the source code of the H.323 Beacon. We thank Christine Dorsey and Susan Evett for their suggestions on an earlier draft of this paper. We also thank Bob Dixon, Pankaj Shah, members of the Internet2 End-to-End Performance Initiative and many others who have provided us valuable feedback at various stages of tool development.

8. REFERENCES

[1] NetIQ's chariot. http://www.netiq.com/products/chr/default.asp.

[2] OARnet H.323 Beacon: A tool to troubleshoot end-to-end H.323 application performance problems. http://www.itecohio.org/beacon.

[3] OpenH323 Resources. http://www.openh323.org.

[4] Spirent SmartBits network measurement suite. http://www.spirentcom.com.

[5] Megaconferences: World's largest Internet Conferences. http://www.megaconference.org/.

[6] Mmdump developed at AT&T research labs. http://www.research.att.com/info/Projects/mmdump.

[7] ITU-T Recommendation H.323, infrastructure of audiovisual services- systems and terminal equipment for audiovisual services, Series H: Audiovisual and multimedia systems, 1999.

[8] P. Calyam, M. Sridharan, W. Mandrawa, and P. Schopis. Performance measurement and analysis of H.323 Traffic, PAM Workshop, April 2004.

[9] L. Lustosa, L. Carvalho, P. Rodrigues, and E. Mota. E-model utilization for speech quality evaluation over VoIP based communication systems. 2003.

[10] H. Schulzrinne, S. Casner, R. Frederick, and V. Jacobson. RTP: A transport protocol for real-time applications. *RFC 1889*, 1996.

[11] A. Tirumala, F. Qin, J. Dugan, J. Ferguson, and K. Gibbs. NLANR Iperf version 1.7. http://dast.nlanr.net/Projects/Iperf.