OSCnet Measurement Project

Report of the combined projects:

Third Frontier Network Collaborations

And

Cincinnati Education Loop Education Testbed

OARnet, Cincinnati State University, Kent State University, The Ohio State University, Southern State Community College, University of Cincinnati, University of Toledo

Project Director:
Steven I. Gordon
Ohio Supercomputer Center
1224 Kinnear Road
Columbus, OH 43214
614-292-4132
sgordon@osc.edu

Report by:
Steven I. Gordon, OSC
Prasad Calyam, OSC

1. Introduction

The OSCnet Measurement project is focused on the building of a testing and evaluation infrastructure for the emerging OSCnet backbone. In particular, we are preparing for the broader use of voice and video technologies, the transfer of large datasets, and the use of other high-performance applications that will make extensive use of network bandwidth. The main goals of the project are as follows:

1) Identify end to end performance bottlenecks in the OSCnet backbone by building a comprehensive *Network Measurement Infrastructure (NMI)*.
2) *Test new and advanced technologies* (H.323/SIP based Voice and Videoconferencing, MPEG3, HDTV, Multicast, Bulk FTP) and equipment (Video streaming Caches, Firewalls, Intrusion Detection Systems, Traffic shapers) before wide-scale adoption in the OSCnet Higher Education communities.
3) *Define the standards for implementation of new technologies in campus networks* so that video and other services run without problems or conflicts with other network devices.
4) *Bring awareness and also train campus-networking professionals* to make optimum use of the capabilities of the OSCnet, so their campus network infrastructure can be upgraded suitably.

Thus, by the end of this project, we hope to have in place a testing infrastructure, proven testing methodologies, and relevant training material that can be employed by all the Ohio-institutions that are planning on adopting and investing in large-scale use of high-performance applications on the OSCnet.

Initially, each of the participating partners acquired upgraded network devices to be used in testing as the OSCnet came online and to serve as production devices for their useful life. These included new IP video equipment, higher capacity firewalls and packet shapers, caching servers, and several ancillary video devices. We also acquired servers to be used in a measurement infrastructure that is to become a permanent part of the OSCnet for monitoring and troubleshooting network problems.

With the equipment in place, the remainder of our efforts have centered around the creation of a set of tools to monitor and analyze the quality of the network end-to-end critical to understanding how, why, and where network problems are reducing the quality of video, audio, and data transfer services. The measurement infrastructure has led to some ideas on how to solve potential network problems through changes in campus network design as well as through the reconfiguration of devices and addition of new approaches. As the testing of those actions are more thoroughly tested over the next few months, we will disseminate that information to all of higher education in the state and assist campuses with their implementation. The remainder of this report summarizes our accomplishments to date and the plans for the future of the project.

2. Project Accomplishments

2.1 Upgrades to Network Appliances

The most direct benefit of the project was the ability of the participating institutions to upgrade their network appliances in preparation for the startup of the Third Frontier Network. Upgraded video equipment meant that the campuses could immediately take advantage of the availability of increased bandwidth to use higher quality video conferencing in their operations. The upgrading of firewalls and packet-shapers so that they could handle the increased bandwidth without slowing the flow of traffic was also important. Finally, the project allowed the technical staff at all of the institutions to gain a deeper understanding of the hardware and related software and how to manage those resources on the more robust OSCnet.

2.2 Network Measurement Infrastructure Development and Deployment for OSCnet

We have developed a number of measurement tools to assess the status of the network as people troubleshoot problems with network performance. These have been developed as open-source tools that are available to the entire community.

2.2.1 H.323 Beacon

We developed the H.323 Beacon [3] tool that can be used to measure, monitor and qualify the performance of an H.323 Videoconference session. It can help an end-user/network engineer/conference-operator, as a debugging tool by providing H.323-protocol specific evidence and other information necessary to troubleshoot Voice and Video over IP (VVoIP) application (e.g. VoIP, Videoconferencing) performance problems in the network and at the host (end-to-end).

The tool is open-source and can run on both Windows and Linux platforms. It has been chosen to be the official troubleshooting tool for the Internet2 Commons (commons.internet2.edu) and the Megaconferences- The Worlds' Largest Internet Videoconferences (www.megaconference.org). The tool has also been quoted in whitepapers of companies such as Polycom and Apparent Networks as a tool they recommend to troubleshoot H.323 Videoconferencing end-to-end performance problems. A company called Visual Networks is working on a re-distribution version of the H.323 Beacon software for its world-wide customer base.

2.2.2 ActiveMon

We developed the ActiveMon [4] tool, which is an extensible and customizable software framework for generation and analysis of active measurements that can be used for routine network health monitoring. It supports various tools that measure network health metrics such as: route changes, delay, jitter, loss, bandwidth and MOS. Using a scalable scheduler called- "OnTimeMeasure", active measurements are initiated in a regulated and non-conflicting fashion between multiple measurement servers distributed at strategic points in a network. It also supports efficient measurements data storage and uses relevant statistical and visualization analyses coupled with alarm generation capabilities, to aid in determining end-to-end network performance bottlenecks along measured network paths.

ActiveMon is currently deployed on OSCnet [5] at several core points in the network backbone and at edges of partner Universities. An alpha software version with detailed instructions on configuring ActiveMon has been developed and is available openly to the University community to deploy on their networks for monitoring their network performance on an on-going basis. In addition, we are discussing potential commercialization and collaboration opportunities of ActiveMon with Plannet Group LLC. It is hoped that ActiveMon will augment their network management product-line and will enable them to better provide managed services relating to network monitoring.

2.2.3 Network Measurement Tools Testing

We compared measurements reported by commercial (E.g. AppareNet, VQMon, Nettach, Hammer, etc.) as well as open-source measurement tools (Ping, OWAMP, Traceroute, Iperf, Pathload, Pathrate, Pathchar, H.323 Beacon, etc.). During the testing, we emulated various end-to-end bottlenecks using WAN emulation and noted whether these tools appropriately identified the different network problems. We also identified and classified tools that were inherently either CPU intensive or Channel intensive. We also studied accuracy of measurements produced due to effects of firewall and packet shapers. Thus, we were able to determine which tools consumed excessive network bandwidth and which devices were friendly for network measurement tools. We appropriately developed measurements scheduling techniques to regulate the bandwidth consumption of measurement tools deployed at strategic points on OSCnet.

2.2.4 Network Measurement Tools Deployment on OSCnet

We deployed ActiveMon in a testbed called "OSCnet Beacon Infrastructure" (TBI), which includes network paths of OARnet OSCnet and Internet2 Abilene, connecting The Ohio State University, University of Cincinnati, University of Toledo, Southern State Community College, Texas A&M University and North Carolina State University. Above testbed monitors network performance between all the participating university partners connected via OSCnet and their performance to Abilene nodes and other national universities connected to Abilene.

Applying advanced data mining techniques on the network measurements data collected from the TBI testbed, we studied issues relating to empirical correlation of network events in a routine monitoring infrastructure. The events corresponded to route changes, bandwidth mis-configuration and anomalies due to measurements mis-reporting. We also used the TBI testbed data to analyze long-term (over a six-month period) network performance trends and quantified end-to-end network performance stability in academic network backbones. In addition, we compared the long-term active measurements with passive measurements (e.g. Netflow, SNMP and syslog) along the TBI testbed paths. By leveraging our experiences in developing and maintaining TBI, we have identified the deployment phases of ActiveMon all across the OSCnet core routers so as to provide routine network performance monitoring for the OARnet Network Engineering team. As a pilot, we have also developed a network "weather-map" [5] that shows both real-time and historic network health data of the TBI testbed paths. Added visualization and analysis functionality are being developed to enhance the data mining of the weather-map.

2.3 Testing Interactive Applications (Ex., audio and video conferencing, remote instrumentation)

2.3.1 VoIP, Videoconferencing and IPTV

We studied the challenges in supporting above interactive audiovisual applications over the Internet. We built a testbed, which included the H.323 Beacon, to collect over 300 traffic traces involving 26 sites in 14 countries, to characterize and model the performance of Voice and Video over IP (VVoIP) traffic (i.e. VoIP, Videoconferencing and IPTV traffic) over the Internet. The sites included participants from: South Korea, Brazil, Poland, Egypt, Ireland, India, China, Australia, Germany, Switzerland, Croatia, Netherlands, Canada, New York, Ohio, Texas, California, North Carolina, Alabama, North Dakota, Pennsylvania, Illinois and Missouri. Our results demonstrated the impact of various voice and video end-point codecs, voice and video packet size distributions, network jitter, lost and re-ordered packets on the end-user perception of audiovisual quality.

Using results of our above studies, we developed a multi-rate VVoIP model framework that reflects the effects of various device factors and network factors that affect the performance of high-speed interactive audiovisual streams, i.e. the ultimate end-user perception of audiovisual quality. We used our VVoIP traffic model in OPNET network simulator studies to evaluate ranges of call-load that can be safely handled in a network path irrespective of other background traffic corresponding to existing network services and end-user applications. We also illustrated how our above framework can be leveraged to perform pre-assessments and on-going assessments to determine if network paths in a VVoIP deployment are capable of supporting high-quality Videoconferencing and IPTV. Above framework can also be used to quantify the effect of video quality degradation in the presence of firewalls, packet shapers, NATs and any other intermediate network devices that affect VVoIP applications.

2.3.2 Video Streaming and Caching

We developed detailed test cases and evaluated video streaming performance in a network via firewall and NAT device connections. We also performed measurements using a caching server as an intermediate network device. The video streaming throughput measurements involved transferring various multimedia streaming data with different file samples (ex., Matrix Reloaded clip, Soldier_300k clip, etc.) using various media players like Windows Media Player, Real Player and Quick time Player. The results produced media stream performance characteristics such as: bandwidth, frame rate, frame rate, packet size, RTT, frames skipped, packet rate, etc.

We studied the impact of file length, encoding rate, encoding protocol, bandwidth restrictions, caching server presence, live vs. on-demand streaming on various measurement parameters. Further, we studied the effects of multiple simultaneous downloads and consecutive downloads. The response time for these tests gave us the measure of benefit obtained by using a caching server. The results demonstrated the maximum achievable throughput and also helped in identifying any end-to-end network performance bottlenecks caused due to intermediate network devices (e.g. firewalls, packet shapers) that hinder obtaining maximum video streaming quality achievable on network paths. We collaborated with Southern State Community College who were provided with a Caching server to make effective use of much needed bandwidth on their bandwidth-constrained network connection to OSCnet.

2.3.3 File transfer Applications (Ex., mass data transfers to remote databases and computing sites)

We studied the networking and application tuning issues that need to be addressed and managed during bulk file transfer applications. We performed experiments in a LAN and on a pilot testbed involving OSC's mass storage site and Wright State University with specific bulk file transfer needs. Using automated scripts in our experiments, we studied experiment scenarios involving multiple TCP Window Sizes (RTT*BW± 100), specialized TCP stacks (Reno, Fast TCP, HS TCP, BIC TCP) with various popular Bulk FTP applications (FTP, SCP, BBFTP, Iperf) and with different application buffer/file sizes. The experimental results indicated the limiting factors that hinder large-scale data transfers over the Internet. In addition, our results demonstrated that adequate network tuning and management could significantly improve the efficiency of bulk data transfers for applications involving data backup and retrieval operations.

For above experiments, we developed automated measurement scripts and well-defined guidelines that effectively help to determine the appropriate system parameters to be tuned for bulk file transfers on any chosen end-to-end WAN path. The automated tools used an iterative approach using different combinations of system tuning parameters based upon accurate bandwidth-delay product measurements on a host-to-host or end-to-end network path. Our results were used by OSC and Wright State University in the "Remote Backup of Critical Administrative Data for Disaster Recovery" project.

2.4 Training and Outreach

The project plans, specific application test plans and their results were periodically discussed with the project partners and posted on a website [1]. A mailing list also was setup for project discussions. The results of the testing activities were compiled in the form of network case studies and project experiences. All of these were disseminated to the community via research papers, technical reports, training workshops at OARnet for university network engineers. As part of our training and outreach efforts, we even had site visits with partner university teams for one-on-one discussions of test plans, measurement equipment installation/reconfiguration and to discuss their progress in relation to the overall project goals. At the end of the report is a list of related publications and presentations, which can be downloaded at the OSC's Networking Research Group's website [2].

2.4.1 Project Publications (Novel Techniques Developed/Experiences/Case Studies):

- P. Calyam, C. G. Lee, A. Sukhov, L. Jorgenson, "A "GAP-Model" for Monitoring Perceptual Quality of VVoIP Applications", Under Review, 2006.
- P. Calyam, C.G.Lee, P. K. Arava, D. Krymskiy, "Conflict-free Scheduling Algorithms for Orchestrating Network-wide Active Measurements", Under Review at IEEE Transactions on Computers Journal, 2006.
- P. Calyam, P. K. Arava, C. Butler, J. Jones, "Network Tuning and Monitoring for Disaster Recovery Data Backup and Retrieval", Under Review at IEEE LCN Network Measurements Workshop, 2006.
- K. Salah, P. Calyam, M. I. Buhari, "Assessing Readiness of IP Networks to Support Desktop Videoconferencing using OPNET", Submitted to Computer Communications Journal, 2006.
- P. Calyam, C.G.Lee, P. K. Arava, D. Krymskiy, "Enhanced EDF Scheduling Algorithms for Orchestrating Network-wide Active Measurements", IEEE Real-Time Systems Symposium (RTSS), 2005.
- P. Calyam, D. Krymskiy, M. Sridharan, P. Schopis, "Active and Passive Measurements on Campus, Regional and National Network Backbone Paths", IEEE International Conference on Computer Communications and Networks (ICCCN), 2005.
- P. Calyam, P. K. Arava, C. Butler, J. Jones, "Bulk-data Transfer Throughput Measurements for Disaster Recovery Data Backup", Technical Report of OARnet/OSC presented at SUN's SuperG Workshop, 2005.
- P. Calyam, S. Bhagavat, P. K. Arava, "Video Streaming Cache Performance Testing for Windows Media, Real Media, QuickTime Media", Technical Report of OARnet/OSC, 2005.
- P. Calyam, C.G. Lee, "Characterizing voice and video traffic behavior on the Internet", International Symposium on Computer and Information Sciences (ISCIS), Proceedings published by Imperial College Press in a special edition of "Advances in Computer Science and Engineering" Book Series, 2005.

- A. Sukhov, P. Calyam, W. Daly, A. Illin, “Towards an Analytical Model for characterizing behavior of High-Speed VVoIP applications", TERENA Networking Conference (TNC), Proceedings published by PSNC in “Journal of Computational Methods”, 2005.
- P. Calyam, C.G. Lee, P.K. Arava, D. Krymskiy, D. Lee, “OnTimeMeasure: A Scalable Framework for scheduling active measurements”, IEEE/IFIP End-to-End Monitoring Workshop (E2EMON), 2005.
- P. Calyam, D. Krymskiy, M. Sridharan, P. Schopis, "TBI: End-to-End Network Performance Measurement Testbed for Empirical-bottleneck Detection", IEEE TRIDENTCOM, 2005.
- A. Sukhov, P. Calyam, W. Daly, A. Illin, "Network Requirements for High-Speed Real-Time Multimedia Data Streams", IPv6 Global Summit, 2004.
- P. Calyam, W. Mandrawa, M. Sridharan, A. Khan, P. Schopis, “H.323 Beacon: An H.323 application related end-to-end performance troubleshooting tool”, ACM SIGCOMM Workshop on Network Troubleshooting (NetTs), 2004.
- P. Calyam, M. Sridharan, W. Mandrawa, P. Schopis “Performance Measurement and Analysis of H.323 Traffic”, Passive and Active Measurement Workshop (PAM), Proceedings published in Springer Lecture Notes in Computer Science, 2004.

2.4.2 Major Project Presentations/Tutorials/Demonstrations:
- Prasad Calyam, “State-wide Disastery Recovery for Ohio’s Higher Education”, OSC *Windows on the Future*, 2006.
- Prasad Calyam, Paul Schopis, “ActiveMon: Development and Deployment Experiences”, *Joint Techs*, Albuquerque, February 2006.
- Prasad Calyam, “Third Frontier Network Monitoring”, *Supercomputing*, Seattle, November 2005.
- *OARnet/Internet2 Network Performance Workshop*, September 2005.
- Prasad Calyam, Paul Schopis, “Tutorial: Troubleshooting Voice and Video Performance over IP using Measurement Tools”, *Measurement Tools Tutorial at Joint Techs*, Salt Lake City, February 2005.
- Prasad Calyam, "OSCnet Measurement Project", *OARTech*, Columbus, December 2004.
- Prasad Calyam, Mark Cain, "OSCnet Measurement Project", *Joint Techs*, Columbus, July 2004.

3. Project Expenditures

Table 1 shows the expenditures of funds by major category. To date, we have spent $581,177 on equipment and software. This included a mix of video hardware, caching servers, firewalls and packet shapers and servers. The equipment is all in use at the partner institutions. OARnet has retained several of the firewall and video components in use for a continuing testbed. Several commercial network measurement packages were also acquired and used in the testing as discussed above.

An additional $177,000 was expended on programming efforts relating to the measurement testbed and the testing program. That has put us in a position to deploy the testbed to the full network and to complete some additional work to advise all campuses on the design of their networks to achieve acceptable performance for video and related services.

Table 1: Project Expenditures by Category

Hardware and Software	
Video	$76,240
Caching	$90,720
Software	$179,525
Firewall	$173,388
Servers	$56,857
Other	$4,447
Subtotal	**$581,177**
Programming work	$176,876
Total Expended to Date	**$758,053**
Planned Expenditures	
Video and Storage Testbed	$170,467
Programming work	$101,468
Grand Total	**$1,029,988**

The remainder of the funds are planned to complete two major efforts. First, we will fully deploy the measurements infrastructure to the entire OSCnet. Measurement servers will be deployed at every PoP (point of presence) of the network. We will also train local administrators how to use the measurement infrastructure to troubleshoot network problems on their own campuses. We will encourage them to add additional measurement servers inside their campus boundaries to aid in this effort.

The second effort, discussed further below, will use current open source gatekeeper technology to create an easily deployable video/Voice over IP gateway service that resolves the problems we found with campus firewalls. Those projects are expected to be completed within the next nine months.

4. Future Work: Resolving Problems with Campus Firewalls

One of the critical findings of our work is the impact of campus firewalls on video and audio services. Resolving these problems will require two sets of actions:

- Deploying and using network measurement tools to troubleshoot network problems end-to-end
- Changing the campus network management to place video devices outside firewalls in a network DMZ or deploying services that allow video and other acceptable traffic to bypass the firewall

We will undertake both tasks to complete our work on this grant. First, we will fully deploy the network measurement infrastructure to the entire OSCnet and train campus network managers how to utilize those tools to troubleshoot network problems. As the infrastructure is deployed, we will create easy-to-use install packages and encourage campus network administrators to add their own, on campus servers to the measurement network. This will allow OARnet to directly assist campuses as they deploy video and related network services and face problems with end-to-end network quality.

Second, we will design and test solutions to typical network video problems. There have been solutions that recommend isolating the H.323 video equipment using a De-Militarized Zone (DMZ). The DMZ by-passes the campus network firewall/NAT and thus avoids performance bottlenecks for video traffic from the H.323 equipment. However, such an approach does not scale to campus-wide end-users, who have to go to DMZ locations to perform routine H.323 videoconferences. Although, vendors such as Radvision have proxy solutions such as "Radvision ECS Gatekeeper/Pathfinder" to let H.323 video streams through firewalls/NATs from any campus location, these solutions are proprietary and are not interoperable with other vendor H.323 equipment.

Recently, there have been success stories [7-8] in the Internet2 and ViDeNet H.323 communities, which have used the open-source "GNU Gatekeeper Proxy" [6] (GNU GK) to overcome the above H.323 firewall/NAT traversal problem. By opening the firewalls for the GNU GK traffic, which alone is in a DMZ, all the H.323 traffic registered with GNU GK is tunneled via the firewall to any remote site without any blockage. Note that in this case, the H.323 equipment need not be located within the DMZ. The goals of this part of the project are to test the above GNU GK solution, validate its utility for firewall/NAT traversal of H.323 video traffic, and assist in its deployment throughout the OSCnet.

The issues, test results, and deployment techniques will be presented in relevant seminars and workshops in Ohio for audiences that include campus network engineers, CIOs and members of industry and academia advisory committees. The test results will also published in the form of a whitepaper, which will include guidelines for Ohio university campuses to setup and maintain the GNU GK solution.

Since this GNU GK software is open source, it has the potential to serve as the gateway for other services and protocols that do not pose a security risk to the network. We will also investigate approaches to authenticate other such services and pass them through the firewall in the same way. The resulting network infrastructure should allow campuses to widely deploy video and other collaboration tools with a minimal risk to network security, allowing the OSCnet to reach its potential as a mechanism for collaboration and cooperation across campuses in Ohio.

5. Conclusions

The project has accomplished many of its original goals and will accomplish some additional goals in the coming months. First, the equipment purchased helped the participating institutions upgrade their network appliances in preparation for the higher bandwidth applications that became available when the OSCnet went into production. The increased capacity along with a deeper understanding of how to manage high bandwidth network resources greatly aided the transition to the OSCnet. Second, the development of the measurement infrastructure will provide a permanent and unique capability to track network problems end-to-end on the network, allowing quick resolution of network problems. That infrastructure will facilitate easier adoption of high bandwidth network applications in the future. Finally, the resolution of the network firewall problem with an easily installed open source server will go a long way toward removing a critical barrier to the regular use of the network for video and other collaborative tools.

References

[1] BOR OSCnet Measurement Project – http://www.osc.edu/oscnet/measurement/
[2] OSC Networking Research Group – http://www.osc.edu/research/networking/
[3] H.323 Beacon Tool – http://www.itecohio.org/beacon
[4] ActiveMon Tool – http://www.itecohio.org/activemon
[5] OSCnet Monitoring Testbed – http://activemon.oar.net
[6] Open-source GNU Gatekeeper Proxy – http://www.gnugk.org
[7] C. Schlatter, "The new ITU Standards for H.323 Firewall and NAT Traversal", SURA/ViDeNet Spring Conference, 2006.
[8] K. Stoeckigt, "Proxy Systems for H.323", SURA/ViDeNet Spring Conference, 2006.