TBI:End-To-End Network Performance Measurement Testbed For Empirical Bottleneck Detection*

[1] Prasad Calyam, [2] Dima Krymskiy, [2] Mukundan Sridharan, [1] Paul Schopis
[1] OARnet, 1224 Kinnear Road,Columbus, Ohio 43212.
Email:{pcalyam, pschopis}@oar.net
[2] Department of Computer Science and Engineering,
The Ohio State University, Columbus, Ohio 43210.
Email:{krymskiy, sridhara}@cse.ohio-state.edu

***Abstract*— Recent advances in networking include new bandwidth-intensive applications, sophisticated protocols that enable real-time data and multimedia delivery and aspects of network security that were not conceived in the beginnings of the Internet. Given these advances and the rapid increase in the number of users accessing the Internet, today's networks need to deliver high levels of end-to-end performance in a reliable fashion. In this paper, we present our novel network measurement methodology which employs an application-specific measurement toolkit including a scaleable test scheduler and analysis module to empirically identify end-to-end bottleneck paths in monitored network routes. To show the utility of our proposed methodology, we present case-studies from a network measurement testbed between 3 University campus labs traversing regional and national academic network backbones. Our case-studies address identifying network measurement anomalies in routine ISP operations due to route changes, device mis-configurations and erroneous data from measurement tools. We also present a performance comparison of campus, regional, national-academic and national-commercial network paths based on the measurement data obtained from our testbed. Finally, we illustrate the requirements and potential of federated measurement testbeds to better characterize end-to-end network performance across multiple ISP domains.**

***Index Terms*— Testbed Cooperation and Integration, Innovative Measurements Methodologies and Tools, Traffic Measurement Testbeds**

I. Introduction

Assessing end-to-end network performance requires performing various kinds of network measurements at strategic points in the network and analyzing the obtained data to make reasonable conclusions on the level of network performance. Since many of the early and basic network protocols did not address various essential requirements for network measurements in their specifications, Network Measurement Infrastructures (NMIs) have become necessary in today's networks. NMIs support network-wide measurement data collection and analyses to detect performance bottlenecks in networks that affect the end-user's experience of network-based applications.

Traditionally simple tools such as Ping and Traceroute were integrated into NMIs [1], [2] to isolate problems relating to network connectivity and network topology. Today's advances in popular Internet applications that use protocols that are complex and bandwidth-intensive, combined with the growing demand for sophisticated network measurements from researchers, ISPs and end-users has led to the need for Next-Generation Network Measurement Infrastructures (NG-NMIs). Some of the notable NG-NMIs being developed in the network measurement community include [3], [4], [5]. The NG-NMIs can be differentiated from the traditional NMIs in terms of:

- A Toolkit that uses a set of more sophisticated tools for active and passive measurements
- A more complex test scheduler that automates and regulates the amount of measurements to be made in the network
- An analysis engine that correlates performance across multiple paths and generates various kinds of alerts; e.g. "Concerned"/"Panic" performance levels
- A publicly accessible measurement infrastructure that permits tests from authenticated end-user machines and generates both host and network related performance data for analysis
- A framework that permits sharing of network measurement data between multiple ISP domains using common measurement data request/response schemas
- An architecture that not only isolates a network bottleneck but also notifies appropriate personnel with information for quick problem resolution

Many of the above mentioned NG-NMIs though attempt to be generic in order to be pertinent in most networks, they fail to scale up as demands and necessities vary vastly in different networks. We are developing a NG-NMI called the TFN-Beacon Infrastructure (TBI) as part of our Third Frontier Network Measurement Project [6] which builds upon the principles of other NG-NMIs but is customized to better suit the measurement and monitoring needs of our Third Frontier Network. The Third Frontier Network (TFN), funded by the Ohio Board of Regents, is the dedicated high-speed fiber-optic OARnet network backbone spanning over 1,600 miles of fiber linking Ohio colleges, universities, K-12 schools and other academic communities together to promote research and

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

economic development. The novel features of our TBI include: an application-specific measurement toolkit, an automated and scaleable test scheduler software and an analysis module to empirically identify end-to-end bottleneck paths and in some cases, even specific bottleneck-hops in network paths being monitored.

We have setup a pilot testbed of measurement-beacons that generate measurement data to analyze network end-to-end performance at the campus, regional and national levels. Our measurement-beacons are located at: The Ohio State University Computer Science Lab (OSUL), University of Cincinnati Computer Science Lab (UOCL), North Carolina State University Computer Science Lab (NCSL) and at the border points within the TFN backbone to which the OSU and UC labs are connected. Our plan is to extend the deployment of the measurement-beacons to many other strategic locations in the TFN based on the experiences and lessons-learned from the pilot testbed. The key research goals of our pilot testbed are:

- To understand network end-to-end performance via partial path and intermediate bottleneck hop analysis,
- To understand network performance measurement data reported by various tools to empirically correlate network events in a routine monitoring infrastructure and
- To compare performance at campus, regional, national-academic and national-commerical network levels to quantify end-to-end network performance stability in the Internet

Towards achieving these goals, we have compiled a few case-studies from the measurement data we have been collecting continuously over a 2-month period between the sites in the pilot testbed. The case-studies address identifying network measurement anomalies in routine ISP operations due to route changes, device mis-configurations and erroneous data from measurement tools. Finally, we present a framework into which we envision the NG-NMIs, such as our TBI, will eventually evolve to meet the future demands of end-users, ISPs and researchers who will seek real-time, historic and forecasted measurement data to better understand network behavior affecting their applications.

The remainder of the paper is organized as follows: In Section II we describe the TBI functionalities. In Section III we describe our testbed and related end-to-end network performance case-studies obtained from our ongoing measurements. In Section IV we discuss related-work and our vision of a futuristic framework for NG-NMIs. In Section V we conclude our work.

II. Measurement Infrastructure

The first contribution of this research study is the development of the software necessary for the TBI to generate end-to-end performance data. The TBI consists of the basic components expected to be seen in any NG-NMI. Fig. 1 shows the various components present in a typical NG-NMI. The TBI components are:

- A set of measurement servers equipped with a measurement toolkit and located at strategic points in the network being monitored

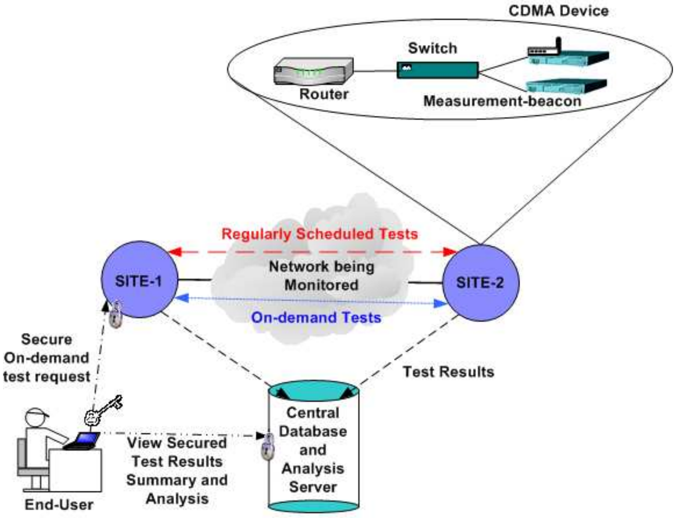
CDMA Device
Switch
Router
Measurement-beacon
Regularly Scheduled Tests
SITE-1
Network being Monitored
SITE-2
On-demand Tests
Secure On-demand test request
Test Results
Central Database and Analysis Server
View Secured Test Results Summary and Analysis
End-User

Fig. 1. Typical NG-NMI Setup

Measured Network Characteristics	Tool	Availability
Round-trip delay	Ping	Open Source
High-precision one-way delay	OWAMP	Open Source
Topology and route changes	Traceroute	Open Source
Bandwidth Capacity: Per-hop	Pathchar	Open Source
Available bandwidth	Pathload	Open Source
Bottleneck bandwidth	Pathrate	Open Source
TCP throughput	Iperf	Open Source
Performance of H.323 audio/video sessions	H.323 Beacon	Open Source
Performance of multicast sessions	Multicast Beacon	Open Source
End-host performance	NDT	Open Source
Network traffic quantity and description	Netflow	Open source
Hop-by-hop path performance	appareNet	Commercial
Application response time	NetQoS	Commercial

TABLE I
Tools used in the Network Measurement Toolkit for TBI

- A central server that launches the tests scheduler and also hosts a database to centrally collect and analyze the data
- Security and trust policies that prevent abuse of the infrastructure

The analysis component of the central server also includes a data visualization and alarm reporting component to notify appropriate personnel who may need to resolve the problem that caused the alarm. The various novelties of our TBI are discussed in the following subsections.

A. Application-specific Measurement Toolkit

It is known that different applications have disparate performance expectations from the network. E.g. large-scale data transfers require high throughput and low loss while voice/video applications demand low latency, jitter and loss. To verify network performance, we need to analyze measurement scenarios where actual application traffic performance is measured. Common methods of verifying performance of a stream of ICMP or UDP packets (Ping, Iperf, etc.) may not necessarily represent performance bottlenecks experienced by actual user-application's traffic. Hence application-specific measurements combined with traditional ICMP and UDP packet stream measurements could provide better insight of the end-to-end network performance experienced by the user-applications. The network measurement toolkit used in the TBI, shown in Table 1, measures end-to-end network performance for web server and database server response times. It also measures end-to-end performance of high-performance applications such as real-time multimedia delivery (e.g. Videoconferencing), multicast applications and large-scale data transfers.

After vigorous acceptance testing of various tools in a LAN environment with WAN emulation, a comprehensive toolkit was chosen. The acceptance testing involved using NISTnet [7] for WAN emulation and Spirent SmartBits [8] for generating cross-traffic. For each test scenario, the testbed switch ports and counters were monitored to determine the type and amount of traffic generated by the various tools. The goal of the acceptance testing was to identify tools that covered multiple measurement metrics, tools that were non-intrusive in nature and tools that responded to the various WAN emulation scenarios in a reasonably accurate manner. The final network measurement toolkit shown in Table I consists of both active and passive techniques and open source solutions as well as commercial solutions. One of the open source tools used in the toolkit is the H.323 Beacon tool [9] that we have developed.

B. Scheduler

In our LAN acceptance tests, we realized that a few of the measurement tools required dedicated system resources such as: network interface card (NIC), CPU processing, ports and multimedia codecs. One example in this context is that the bandwidth estimation tools such as Iperf require significant amount of NIC resources. Another example is the case of the H.323 Beacon where a test instance requires many ports (e.g. TCP 1720) to be exclusively available. Also, running multiple simultaneous measurements using many tools on a measurement server could affect the results reported by all the tools since the tests interfere with the each other as they compete for the same system and network resources.

Monitoring infrastructures such as Network Weather Service [10] use a token passing algorithm between measurement servers to avoid conflicts in performing simple delay and bandwidth measurements. Token based scheduling methods fail to scale when multiple parallel measurements could simultaneously occur on two links that have no common factors in a network infrastructure involving multiple links. Hence, we have developed a scaleable scheduler that can not only avoid test conflicts but also permits controlling the granularity of the tests and the amount of measurement traffic injected into the network. Fig. 2 shows the workflow of the centralized scheduler that is launched from the 'Central Database and Analysis Beacon' server which is used to centrally collect all the measurement data. The central server also has the capability to analyze, summarize and visualize measurement data in real-time.

The 'Central Scheduler Module' is responsible for generating and updating Timetables in real-time for each measurement server in the testbed. The timetables specify the times at which cron jobs are executed for launching measurement tests on the measurement-beacons. The central scheduler uses two types of data as input for the timetable generation algorithm: topology description and measurement tool description. The topology description data consists of:

- List of sites to be monitored and nodes-per-site
- Sites in full-mesh, tree or hybrid-mesh

and the measurement tool description data consists of:

- Tool name
- Measurement characteristic
- Tool associated command with relevant options
- Tool associated script file
- Uni-directional/Bi-directional measurement data generation
- Estimated test execution time
- Intensity of impact (low/high)

The scheduler uses the topology information to determine the 'Randomized List of all Tests'. The randomization of the list of tests between sites prevents a fixed periodicity in measurements and thus enables capturing events that occur outside the fixed periods of scheduled measurements. The tests in the lists are integrated into the appropriate timetables in a round-robin fashion by the scheduler. Figs. 3, 4, 5 show the full-mesh, tree and hybrid-mesh topologies respectively. A full-mesh topology involves testing from each measurement site to every other measurement site, a tree topology involves testing between neighboring sites only and a hybrid-mesh topology involves supporting subsets of servers that need a combination of full-mesh and tree topology type measurements.

The scheduler aims at bidirectional testing of the paths between measurement servers in order to explore effects of the asymmetries in the Internet pathways on the relative performance between any two measurement sites. It takes advantage of uni-directional or bi-directional data generation supported in measurement tools. E.g. OWAMP or H.323 Beacon present data from measurements on both ends of the test for one test execution instance. Hence one test is sufficient to obtain measurement data in both directions of a path.

The scheduler sorts the 'Tool Data List' to permit customization of tests at any desired frequency and quantity. The sorting is based on the tool specific factors such as duration of the test, execution time and the low/high intensity of impact of the tool on the network. Low impact tools refer to the light-weight tools such as Ping, OWAMP and Traceroute and the high impact tools refer to the tools that inject relatively larger number of packets and take longer periods of time to complete. Iperf, H.323 Beacon and pathrate are examples of high-impact tools. The timetables generated are updated in real-time on the measurement servers to orchestrate network-wide measurements without overlapping of tests at any point in the network.

An open-issue with regard to the scheduler which we are currently investigating addresses integrating on-demand tests into the measurement schedules of the measurement-beacons.

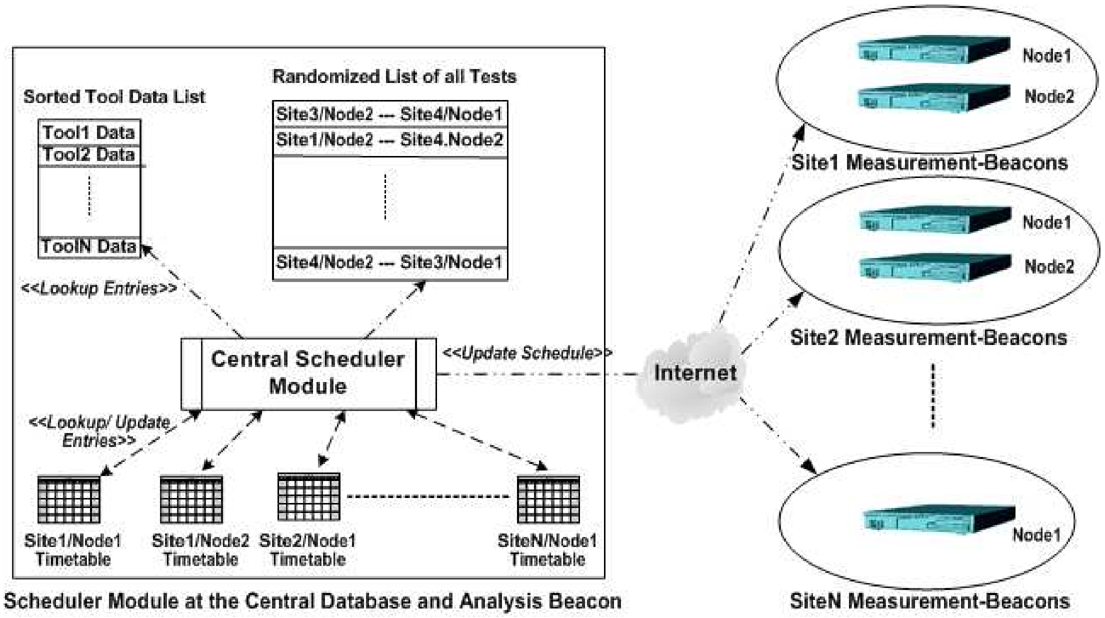
Sorted Tool Data List
Tool1 Data
Tool2 Data
ToolN Data
Randomized List of all Tests
Site3/Node2 --- Site4/Node1
Site1/Node2 --- Site4.Node2
Site4/Node2 --- Site3/Node1
<<Lookup Entries>>
Central Scheduler Module
<<Update Schedule>>
Internet
<<Lookup/ Update Entries>>
Site1/Node1 Timetable
Site1/Node2 Timetable
Site2/Node1 Timetable
SiteN/Node1 Timetable
Scheduler Module at the Central Database and Analysis Beacon
Node1
Node2
Site1 Measurement-Beacons
Node1
Node2
Site2 Measurement-Beacons
Node1
SiteN Measurement-Beacons

Fig. 2. Network-wide Measurement Tests Scheduler Workflow.

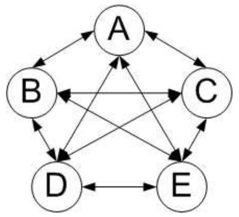

Fig. 3. Full Mesh Topology.

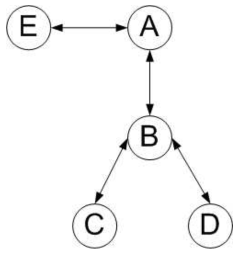

Fig. 4. Tree Topology.

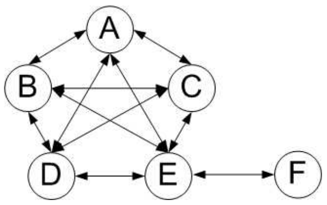

Fig. 5. Hybrid-Mesh Topology.

In our current implementation, for regularly scheduled tests, we specify a start time and an end-time value of ∞. A value of ∞ implies the tests are expected to occur in a recurrent fashion. For the on-demand tests, we are developing an approach in which an on-demand test will be treated as a specialized test with start and non-infinite end times specified for each test in the measurement tool description. When an on-demand test has to be included in the system, a revision of the timetables and a live update of the timetables on all the measurement-beacons will be performed by the 'Central Scheduler Module'.

C. Analyzer

The analysis functionality at the central database forms the most critical and challenging component of our TBI. We have developed a basic version of our analyzer and are currently enhancing its feature set. Our main design goals for developing the various analysis functionality are:

- To effectively identify anomalies (with a low probability of false-alarms) that result in a decrease of the network end-to-end performance
- To perform multi-path data correlation to isolate performance problems involving multiple-links
- To obtain a network-wide performance view via a "weather-map" functionality

The alarms generated by the analyzer are based on the 'Good', 'Acceptable' and 'Poor' performance thresholds we pre-define for the various measurement metrics such as delay, bandwidth, loss and jitter. The thresholds are set based on application-specific performance considerations since we are interested in understanding the user-experience of the network for real-time and data intensive applications. Recently, we developed network performance thresholds for voice and video conferencing applications over the Internet [11] for good, acceptable and poor audiovisual quality grades as perceived by an end-user. We are extending this methodology for determining performance thresholds for other applications such as multicast applications and large-scale data transfers. We are also investigating the possibility of determining a stability metric for the monitored paths based on the measurement data variations between the good, acceptable and poor performance thresholds. Our current version of the analyzer mainly identifies route changes that affect end-to-end performance of any of the links and tracks cases where there are decreases in the performance over links due to route changes and other anomalies that affect end-user network-based applications. An extensive visualization component has also been built into the analyzer that instantly generates varies perspectives of the data collected over the multiple links. The schema used for storing measurement data in our MySQL database enables the analyzer to perform data mining in an effective manner to create network-wide views. These views enable us to identify not only bottleneck links but in some cases even bottleneck hops in various paths.

D. Security Issues

The deployed measurement-beacons could potentially be compromised by network hackers and intruders who could abuse the measurement infrastructure to launch Distributed Denial of Service (DDoS) attacks. To prevent such occurrences we enforce security policies and firewall rules (access-lists) both at the measurement-beacons and at the routers that are associated with the measurement-beacon locations. Since IP-spoofing could break these 2 layers of defense, we are also investigating integrating various Access, Authentication and Authorization (AAA) schemes such as shared private keys and MD5 cryptographic checksums used in [2] into our measurement tool scripts in order to further secure our TBI and also to be able to have a trust relationship for measurement tests with other ISPs who run similar measurement infrastructures.

III. Testbed And Data Analysis

A. End-to-End Network Performance Measurement Testbed

Fig. 6 shows the pilot testbed we setup to test the TBI software for regular monitoring of the end-to-end network performance measurement across campus, regional and national-academic network sites. Our measurements currently are directed towards using active measurements only. The frequency of scheduled active measurements in the pilot testbed is once in 2 hours. As shown in the Fig. 6, each measurement site has 2 measurement-beacon servers connected to the routers at strategic points to obtain the end-to-end performance data. CDMA time sources have been deployed at each of the sites, making them Stratum-1 NTP Servers, for precise one-way delay measurements. The Network Time Protocol (NTP) has been configured at all the sites to peer with 4 other neighboring peers to obtain synchronization. In order to obtain the best possible network measurements without any skew from end-host performance issues, the measurement-beacons are dedicated servers with top-of-the-line hardware; i.e. a gigE interface NIC, powerful dual-processor CPU, abundant memory and disk space.

B. Case-studies involving empirical network-anomalies

In the following subsections we describe a few case-studies we collected using the above pilot testbed with the most current version of our TBI software. The case-studies feature data for the anomalies observed in a 2-month monitoring period. We have characterized the anomalies in the case studies based on those caused by route changes in the paths and those caused by network device mis-configurations or erroneous measurements from the measurement tools under certain conditions. The measurement results presented in the case studies also present a good comparison of data generated by various measurement tools. They also illustrate how active measurements from tools respond to the stimulus of changes in the network and report data which could be used to potentially identify end-to-end performance bottlenecks affecting end-user applications.

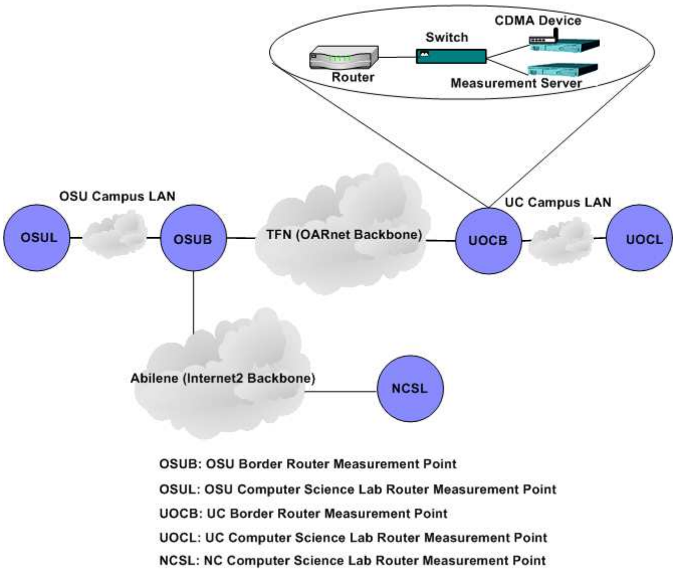
CDMA Device
Switch
Router
Measurement Server
OSU Campus LAN
UC Campus LAN
OSUL
OSUB
TFN (OARnet Backbone)
UOCB
UOCL
Abilene (Internet2 Backbone)
NCSL
OSUB: OSU Border Router Measurement Point
OSUL: OSU Computer Science Lab Router Measurement Point
UOCB: UC Border Router Measurement Point
UOCL: UC Computer Science Lab Router Measurement Point
NCSL: NC Computer Science Lab Router Measurement Point

Fig. 6. Pilot Testbed spanning Campus, Regional and National Academic Networks

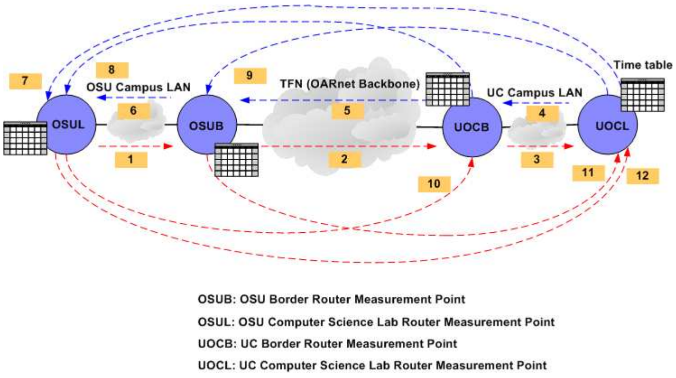
Time table
7
8
9
OSU Campus LAN
TFN (OARnet Backbone)
UC Campus LAN
6
5
4
OSUL
OSUB
UOCB
UOCL
1
2
3
10
11
12
OSUB: OSU Border Router Measurement Point
OSUL: OSU Computer Science Lab Router Measurement Point
UOCB: UC Border Router Measurement Point
UOCL: UC Computer Science Lab Router Measurement Point

Fig. 7. Paths observed for anomalies caused by route changes

1) Anomalies caused by route changes: The most apparent anomalies observed in NMIs are those caused by route changes. Route changes are attributed to route flaps caused by suboptimal routing protocol behavior, network infrastructure failures, re-configuration of networks or load-balancing strategies used by ISPs to improve network performance. They could result in end-to-end performance problems since they contribute to convergence delay of the network and also affect the path properties (round-trip delay, available bandwidth, lost connectivity, etc.)

Fig. 7 shows 12 paths we analyzed in our testbed to identify anomalies due to route changes. The 12 paths arise from measurements performed at each measurement-beacon site to 3 other sites in both directions. Fig. 8 shows our visualization method to represent route changes. We have automated the generation of above plots for route changes seen in routine monitoring. From Fig. 8 we can observe an actual route change that occurred on 15th of July at about 6pm

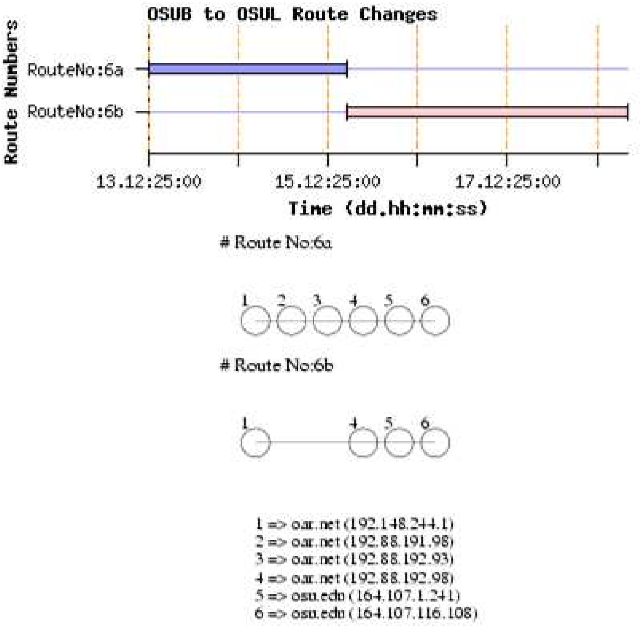
OSUB to OSUL Route Changes
Route Numbers
RouteNo:6a
RouteNo:6b
13.12:25:00
15.12:25:00
17.12:25:00
Time (dd.hh:mm:ss)
# Route No:6a
1 2 3 4 5 6
# Route No:6b
1 4 5 6
1 => oar.net (192.148.244.1)
2 => oar.net (192.88.191.98)
3 => oar.net (192.88.192.93)
4 => oar.net (192.88.192.98)
5 => osu.edu (164.107.1.241)
6 => osu.edu (164.107.116.108)

Fig. 8. Timeline Graph for the Route Change between OSUB and OSUL

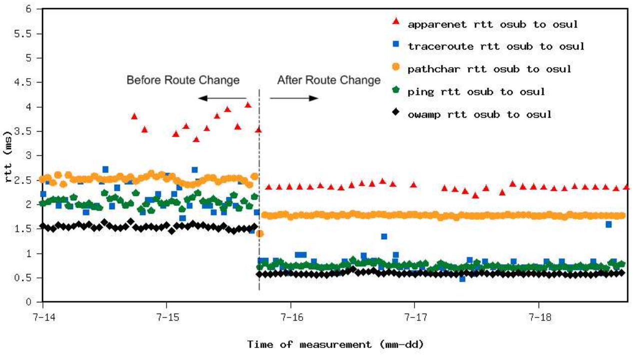

Fig. 9. Network Round Trip Delay Variations between OSUB and OSUL

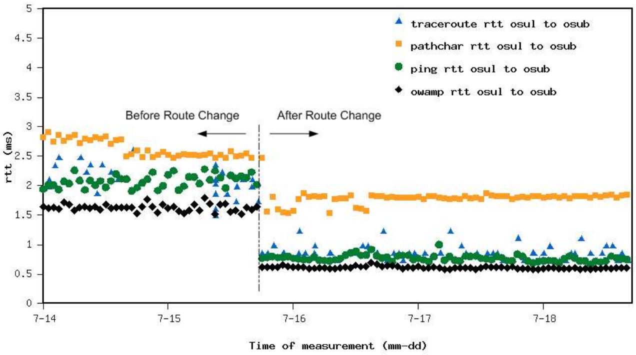

Fig. 10. Network Round Trip Delay Variations between OSUL and OSUB

where a more powerful router was replaced at the OSU border (192.88.192.98) which overrode the previous set of 3 routers (192.148.244.1, 192.88.191.93, 192.88.192.98). Figs. 9 - 11 show the change in the measurements shown by the various tools for this route changes. We can observe an apparent improvement in the performance of the links with the decrease in the round-trip delays and jitter. We also noted an increase in the maximum available bandwidth for both paths (OSUB to OSUL and OSUL to OSUB). The loss reported by appareNet, Iperf and other tools also reduced noticeably. This increased

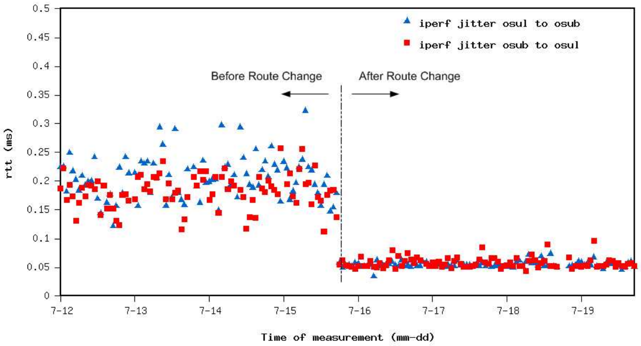

Fig. 11. Network Jitter Variations between OSUL and OSUB

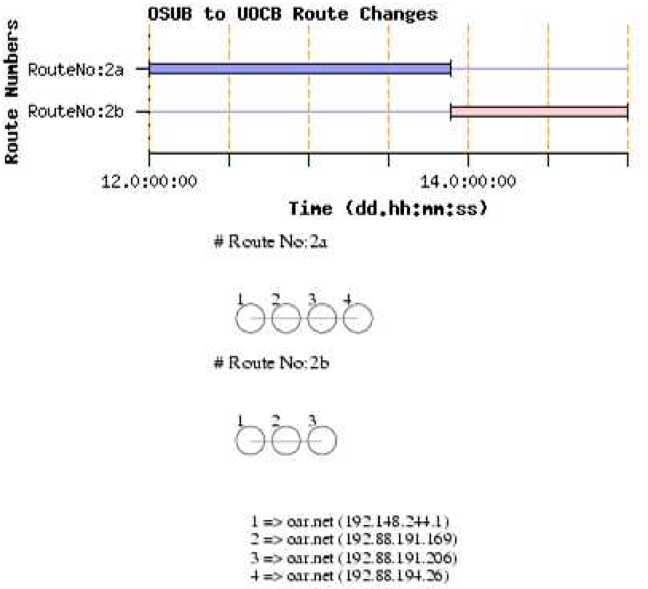
OSUB to UOCB Route Changes
Route Numbers
RouteNo:2a
RouteNo:2b
12.0:00:00
14.0:00:00
Time (dd.hh:mm:ss)
# Route No:2a
1 2 3 4
# Route No:2b
1 2 3
1 => oar.net (192.148.244.1)
2 => oar.net (192.88.191.169)
3 => oar.net (192.88.191.206)
4 => oar.net (192.88.194.26)

Fig. 12. Timeline graph for Route Change between UOCB and OSUB

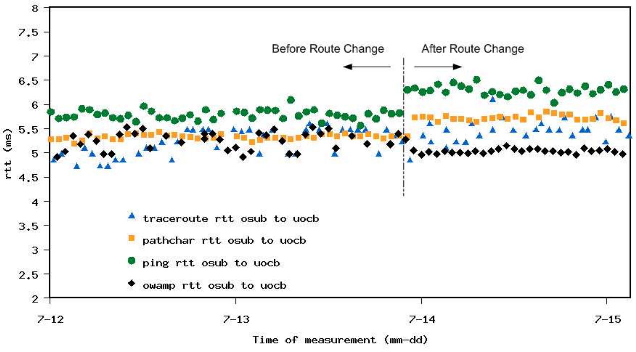

Fig. 13. Network Round Trip Delay Variations between OSUB and UOCB

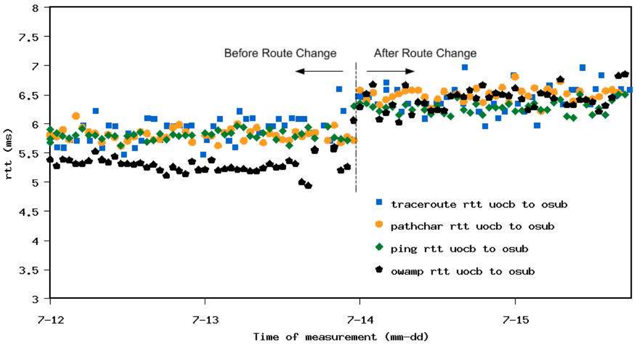

Fig. 14. Network Round Trip Delay Variations between UOCB and OSUB

performance was also noticed in the end-to-end performance measurements between the UC Border measurement-beacon and the OSU Computer Science Lab measurement-beacon. The OSU Border measurement-beacon that experienced the performance problem is located along the path between the UC Border and the OSU Computer Science Lab measurement-beacons.

The results of the above described case were obtained using visual inspection of the measurement data. In an alternate scenario shown in Figs. 13 and 14, the end-to-end path performance actually decreased after a route change (shown in Fig. 12) between OSUB and UOCB. In the cases where such decrease in performance crosses beyond the poor network performance threshold, we have configured the TBI to notify us with an alarm.

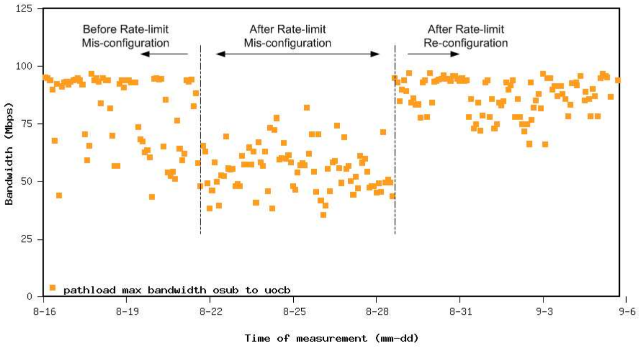

Fig. 15. Pathload variations caused by a device mis-configuration

2) Other anomalies: Besides anomalies caused by route changes, we noticed 2 other significant anomalies from our measurement data. The first anomaly was observed in pathload bandwidth measurements shown in Fig. 15 when a rate-limit was accidentally misconfigured at the router next to OSUB. There was a significant drop in the available bandwidth measurements. The average available bandwidth changed from 80.37 Mbps to 59.85 Mbps in the 100Mbps link. After a few days the misconfiguration was corrected, which resulted in an increase in the average available bandwidth to about 82.53 Mbps for the periods shown in Fig. 15.

The second anomaly was observed in the OWAMP variations shown in Fig. 16. OWAMP tool relies on the NTP clock offset to report one-way delays at the microsecond precision-level. Due to NTP instability in the clocks at UOCB, an interesting pattern of one-ways delays were observed in the measurements involving UOCB. Periodically as the NTP estimates skewed off, NTP compensated for the extreme skew and gave rise to an anomaly we termed as an ”ocean-wave” anomaly of OWAMP.

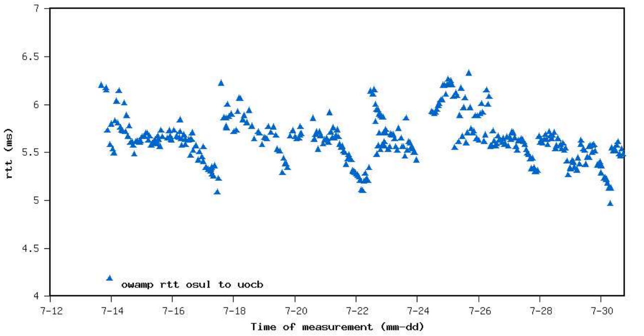

Fig. 16. OWAMP variations caused by NTP behaviour between OSUL to UOCB

C. End-to-End Performance Comparison of Academic and Commercial Networks

In this section, we compare the end-to-end performance of campus, regional, national-academic and national-commercial networks based on the H.323 Beacon objective Mean Opinion Score (MOS) measurements. A MOS measurement for a network link is useful to understand an end-user experience if he/she uses a voice or video over IP application on that particular network link. The objective MOS values are reported on a scale of 1 to 5; 1-3 range being poor, 3-4 range being acceptable and 4-5 range being good. The reader is referred to [9] and [11] for more details pertaining to the H.323 Beacon's objective MOS measurements.

Table II shows the average round-trip delay, average jitter, average number of packets lost, average MOS and average coefficient of variation of MOS, as measured by the H.323 Beacon, for different networks. Campus measurements refer to measurements between 2 campus labs (OSUL and UOCL), regional measurements refer to measurements between 2 regional border routers (OSUB and UOCB), national-academic measurements refer to measurements between 2 national labs (OSUL and NCSL) and national-commercial measurements refer to measurements between a campus site and a commercial site in the New York area. H.323 Beacon tests were run regularly between the above mentioned sites over a period of two weeks and the summarized results have been presented in Table. II. The 'Coefficient of Variation of MOS' values in the table are a measure of the relative variability of the MOS values and are expressed as the standard deviation of the MOS as a percentage of the mean.

As seen from the co-efficient of variation of MOS values in Table II, academic networks are most suitable for voice and video over IP. Our results show, that the average MOS

Site Characteristics	Delay (ms)	Jitter (ms)	Number of Packets Lost	MOS	Coefficient of Variation of MOS (%)
Regional	7	3.167	2	4.4033	0.13
National-academic	45.67	3.667	8	4.3033	0.35
Campus	23.67	5.67	9	4.1733	1.54
National-commercial	46	22	11	3.967	7.4

TABLE II

TABLE SHOWING THE PERFORMANCE OF ACADEMIC AND COMMERCIAL NETWORKS AS MEASURED BY THE H.323 BEACON

values for commercial networks, tend to be lesser and have a greater degree of variance than the measurements in academic networks. Also from Table II we can observe that, amongst the academic networks, the MOS values are most stable in regional network measurements in comparison to the campus and national-academic measurements. The campus measurements are the least stable in comparison to regional and national-academic measurements. The instabilities in the campus measurements can be attributed to LAN performance limitations in department labs. The instabilities in the national measurements can be attributed to the fact that the traffic traverses multiple ISP backbones and has a higher probability of experiencing congestion or other performance bottlenecks compared to the regional network measurements.

IV. NG-NMIS FEDERATION FRAMEWORK

In this section, we describe a framework of a federation of measurement infrastructures that could collectively be used to understand end-to-end performance bottlenecks caused by network issues spanning multiple ISP domains. We envision our TBI to evolve in an inter-operable manner so as to interact with other major network measurement infrastructures globally in a federated fashion.

Since many of the early Internet protocols did not address issues with network measurements in a comprehensive manner, identifying and resolving performance problems in ISP backbones has become a challenging task. Many kinds of active and passive measurement techniques have emerged to address quick detection, identification and resolution of anomalies in the network. However, integrating these techniques into existing devices (e.g. routers and switches) has proven to be non-scaleable since these techniques cause a significant overhead on the routers that could affect their main functionality of routing packets. Since satisfying Service Level Agreements (SLAs) involving high network availability and low thresholds of network performance parameters are critical to remain competitive, ISPs have started instrumenting their networks with NMIs for active measurements (bandwidth, delay, etc.) and passive measurements (Netflow data collection, tracking BGP advertisements, etc.). This has led to the generation of massive amounts of valuable measurement data within individual ISP domains which reflect the status of the network at any given point of time.

Fig. 17 shows a futuristic view of interoperable NG-NMIs in multiple ISP domains where ISPs share measurement data with other trusted ISPs. The reader can refer back to Section 1 where we enlisted the important characteristics of an NG-NMI. The network measurement research community is already addressing issues necessary for developing such a framework. Various techniques for bottleneck identification and quantification of the anomalies across multiple links have been designed [12], [13]. There has been progress even in the area of standards for relevant data schemas [14] such that ISPs could access measurement data of other trusted ISPs in the federation to identify location and cause of the performance bottlenecks affecting end-users in any ISPs network. Some of the open-issues in research that are currently being pursued in the context of NG-NMIs include discovering appropriate measurement-beacons across test paths in real-time and more complex methods of data mining to correlate network wide events to identify mis-behaving flows that disrupt the performance of the traffic of the end-users in the network.

V. CONCLUSION AND FUTURE WORK

In this paper, we presented our novel network measurement methodology we are developing for routine monitoring and end-to-end performance troubleshooting of our Third Frontier Network backbone which connects major academic and research institutions in Ohio. Our measurement methodology includes an application-specific measurement toolkit, a scaleable test scheduler and analysis module to empirically identify end-to-end bottleneck paths in monitored network routes. To show the utility of our proposed methodology, we presented a few case-studies obtained from regular monitoring of a network end-to-end performance measurement testbed between 3 University campus labs traversing regional and national academic network backbones. We presented a performance comparison of campus, regional, national-academic and national-commercial networks based on the measurement data generated by the H.323 Beacon in our pilot testbed. Finally, we illustrated our vision of how our TBI will evolve and interoperate with other major network measurement infrastructures globally in a federated fashion.

We are currently in the process of increasing the number of monitored sites in other universities and institutions that use our network backbone. We are enhancing our TBI software to be capable of more complex techniques of identifying end-to-end performance bottlenecks. We have developed a preliminary version of a network-wide monitoring weather-map which we would like to extend to obtain more interesting visualizations. We are planning to address DDoS attack identification and trace-back mechanisms based on the Netflow data obtained from routers at multiple vantage points in our network. We are also planning to implement various XML-based request/response data schemas being recommended by the Global Grid Forum Network Measurement Working Group [14].

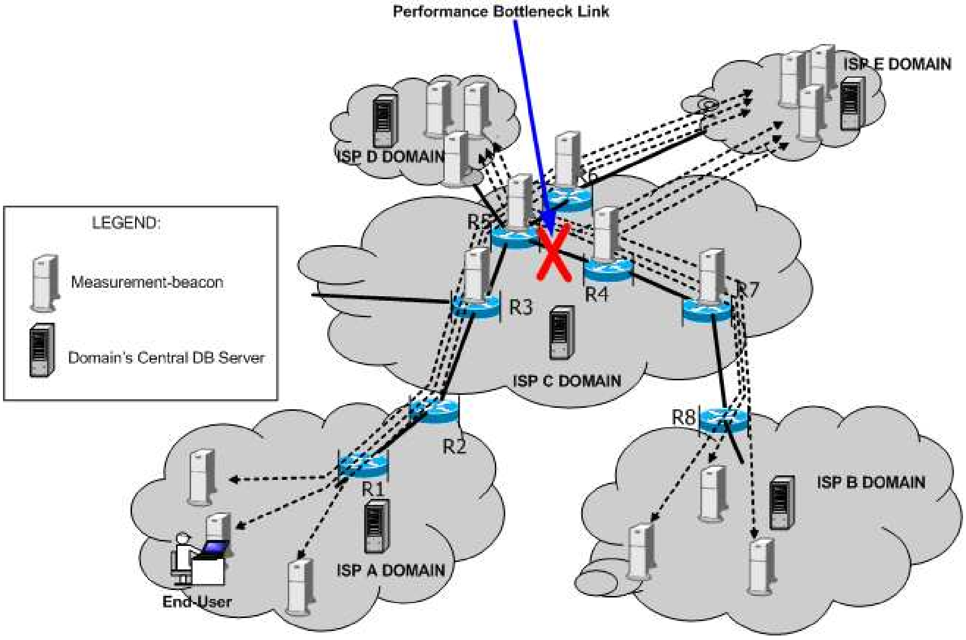

Fig. 17. Futuristic-view of NG-NMIs Federation Framework

REFERENCES

[1] *G. Almes, et. al., "IP Performance Metrics: Metrics, Tools, and Infrastructure, 1997* `http://www.advanced.org/surveyor/`

[2] *V. Paxson, J. Mahdavi, J. Adams, M. Mathis, An Architecture for Large-scale Internet Measurement, IEEE Communications, 1998.*

[3] *L. Cottrell, C. Logg, I. Mei, Experiences and Results from a New High Performance Network and Application Monitoring Toolkit, PAM 2003.*

[4] *T. McGregor, et. al., NLANR Active Measurement Project (AMP), High Performance Networks: Measurement and Analysis Collaborations Workshop, 1999.*

[5] *E. Boyd, et. al., Internet2 End-to-End Performance Initiative Performance Evaluation System, 2003.* `http://e2epi.internet2.edu`

[6] *Third Frontier Network Measurement Project, 2004* `http://tfn.oar.net/measurement`

[7] *NISTnet network emulation package* `http://snad.ncsl.nist.gov/itg/nistnet/`

[8] *Spirent SmartBitsTM Network Measurement Suite* `http://www.spirentcom.com`

[9] *P. Calyam, W. Mandrawa, M. Sridharan, A. Khan, P. Schopis, H.323 Beacon: An H.323 application related end-to-end performance troubleshooting tool, Proceedings of ACM SIGCOMM NetTs 2004* `http://www.itecohio.org/beacon`

[10] *R. Wolski, N. Spring, J. Hayes, The network weather service: A distributed resource performance forecasting service for metacomputing, Proceedings of Future Generation Computer Systems, 1999* `http://nws.cs.ucsb.edu`

[11] *P. Calyam, M. Sridharan, W. Mandrawa, P. Schopis, Performance Measurement and Analysis of H.323 Traffic, Passive and Active Measurement Workshop (PAM 04).*

[12] *A. Lakhina, M. Crovella, C. Diot, "Diagnosing network-wide traffic anomalies", ACM SIGCOMM 2004.*

[13] *P. Barford, J. Kline, D. Plonka, and A. Ron, A signal analysis of network traffic anomalies, ACM SIGCOMM Internet Measurement Workshop, 2002.*

[14] *GGF NMWG Request/Response Schema.* `http://www-didc.lbl.gov/NMWG`