IPCOMM2006
September 26-27 • Gaylord Opryland • Nashville, TN

IPT CASE STUDY

Ohio Colleges and Universities: Adding Video to Your Converged Network

Paul Schopis
Director Network Engineering & Operations/ Associate Director
Third Frontier Network/OARnet

What Will Be Covered

Disclaimer
A little history
The design and building of the TFN
Video and Advanced Services as a demonstration of TFN
Business model
Problems & Solutions

Disclaimer

- I am not associated with Ohio University (my son does go to school there)
- I am associated with OARnet/TFN which is a division of the Ohio Supercomputer Center
- OSC is a project of the State Board of Regents
- OBOR uses Ohio State University as the fiduciary agent for the program e.g. I am an employee of OSU
- OSU and OU are two different institutions
- OSU is a Big Ten School e.g. Football is ***VERY*** important (2002 National Champs)
- OU is a MAC school e.g. Football traditionally is not very good

A Little History

OARnet Background

- Founded in 1987 as part of the Ohio Supercomputing Center
- 90+ higher ed member institutions
- Board of Regents funding
- OSTEER advisory council
- Internet2 GigaPOP

Third Frontier Network

- Phase 1: replace backbone with dark fiber
- Phase 2: connect 17 universities to network with dark fiber or gig circuits
- Phase 3: connect other universities and colleges
- Phase 4: connect other partners

Dark Fiber Acquisition

- RFP issued during Summer of 2002
- Dark fiber was strongly preferred, but leased services considered
- Vendors who bid dark fiber were required to offer a minimum of a single pair of fiber over their network

Dark Fiber Acquisition

- Determined that leased lambdas were too expensive and not widely available
- Selected a bid from Spectrum Networks for single pair of fibers
  - American Electric Power (AEP)
  - Williams Communications (Wiltel)

Dark Fiber Acquisition

- $4.6 M for 20 year IRUs
- $342K/yr for maintenance
- 1600+ route miles
- Truewave, SMF-28, LEAF or Terra Light Fiber
- Aerial and buried

IPCOMM2006

September 25-27 • Gaylord Opryland • Nashville, TN

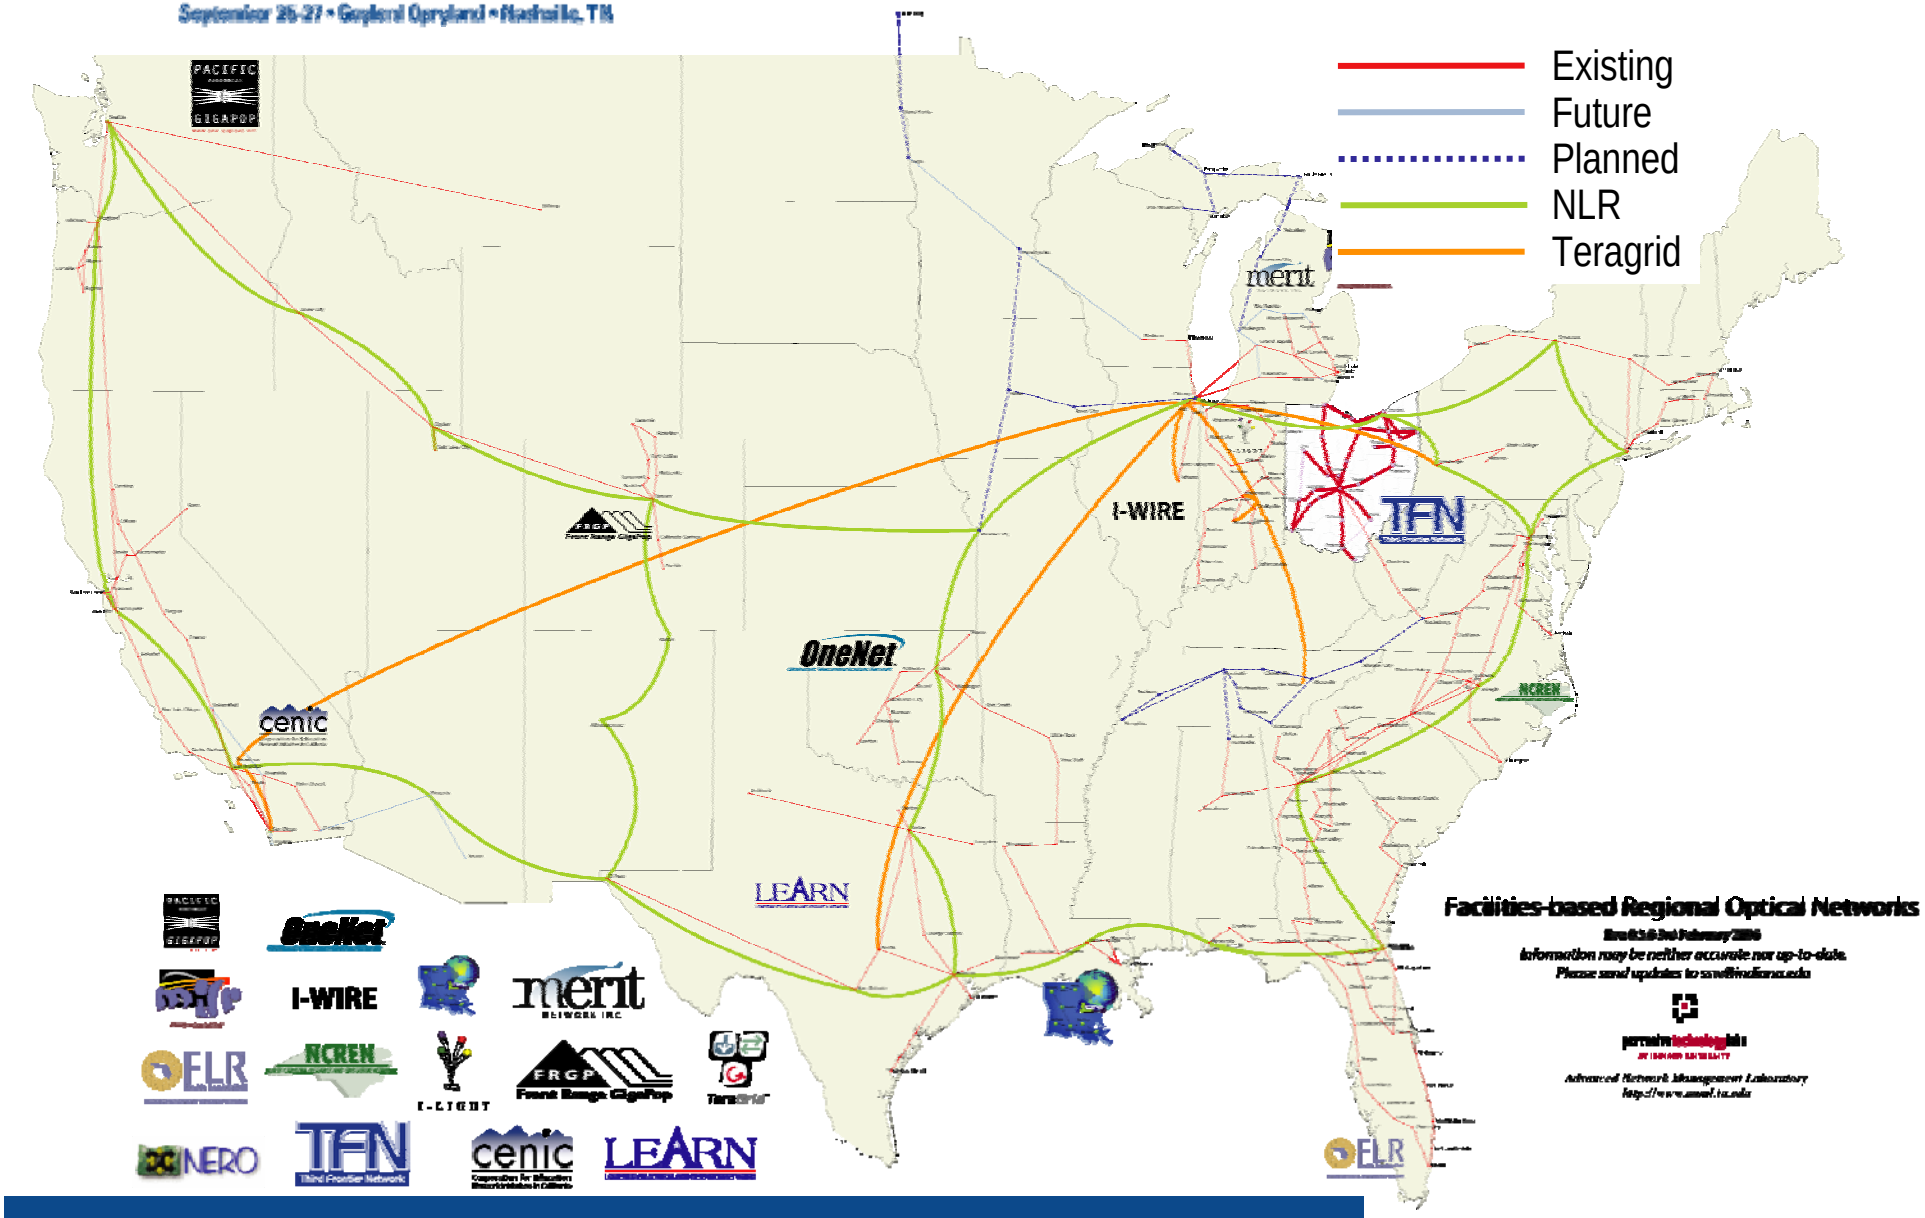

IPCOMM2006

September 25-27 • Gaylord Opryland • Nashville, TN

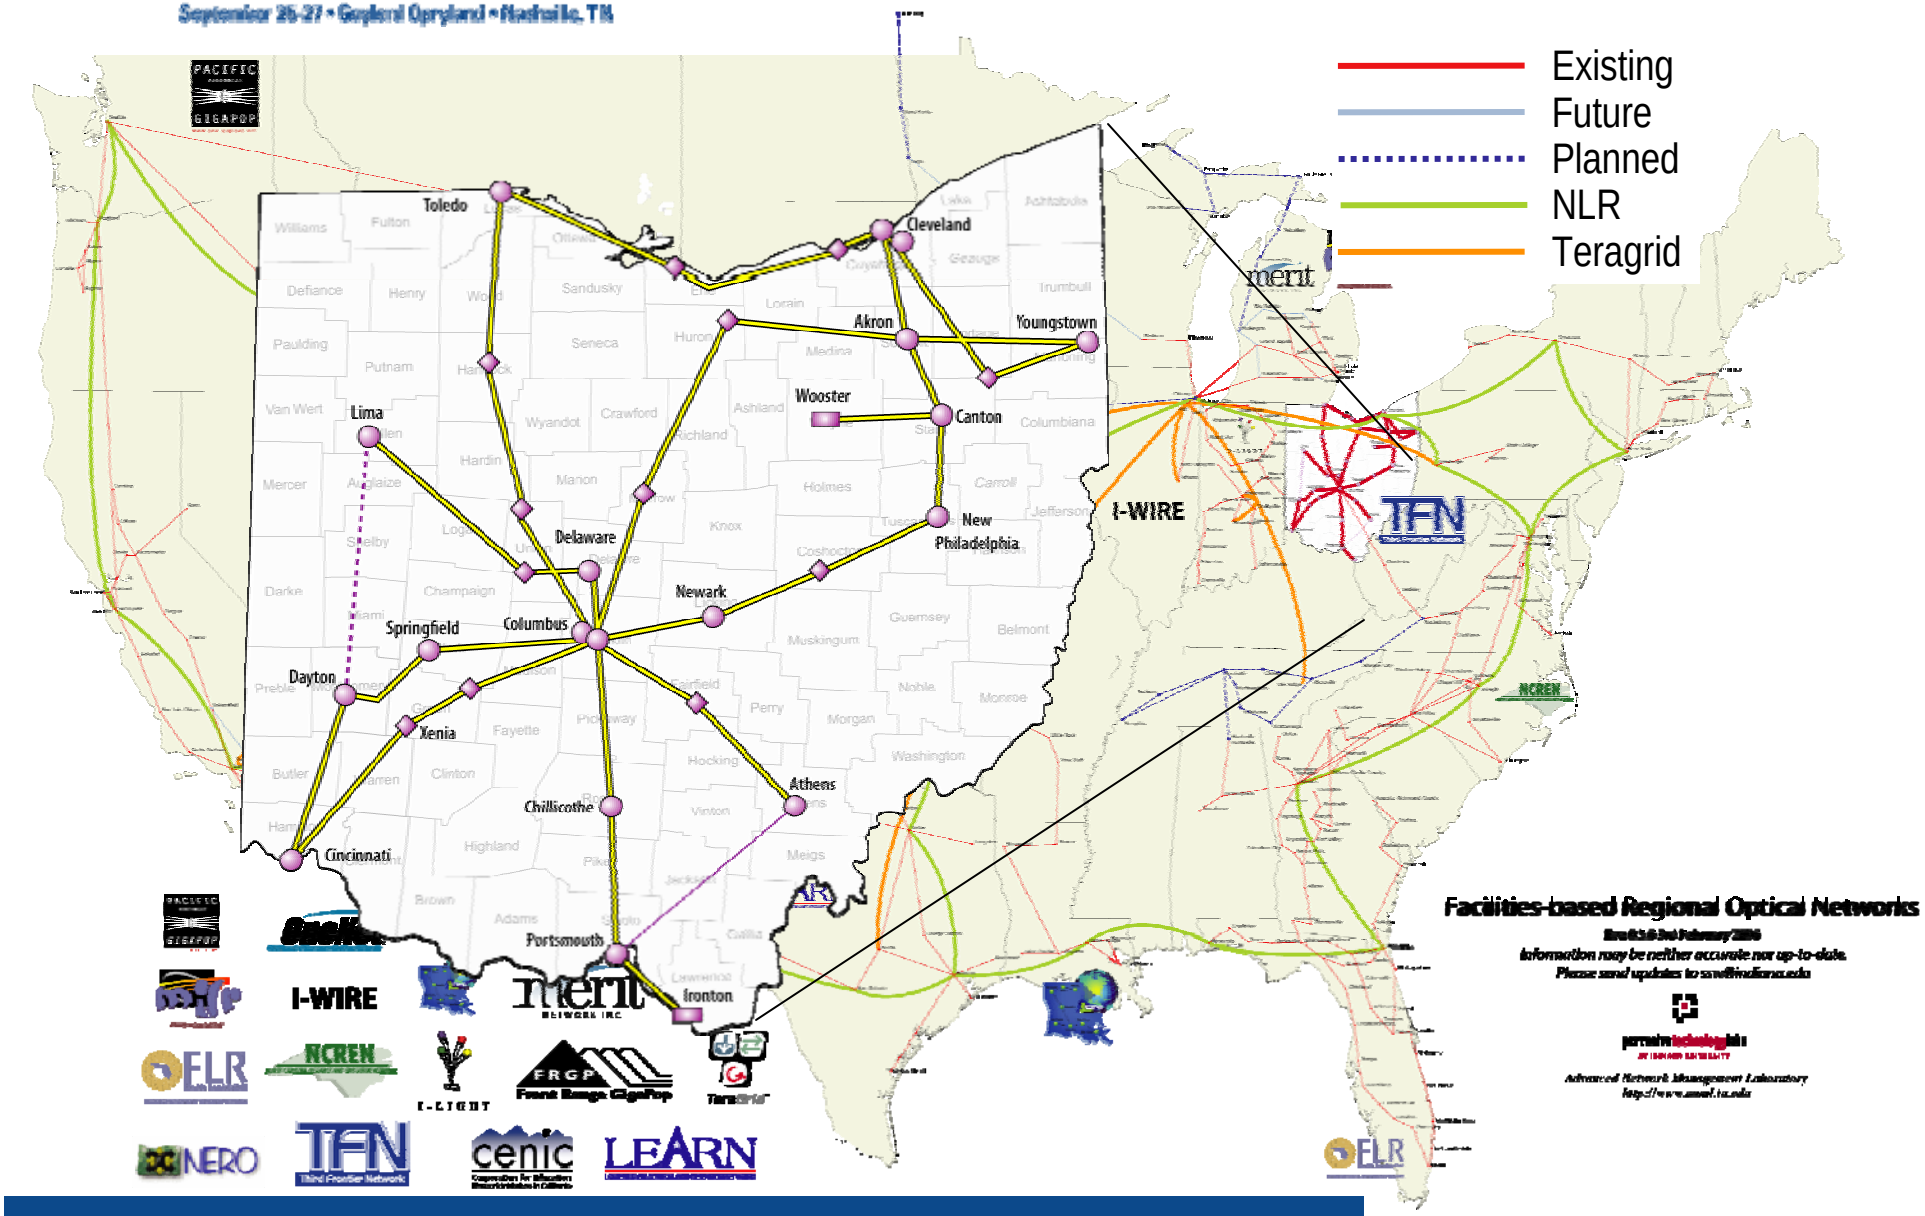


Toledo
Cleveland
Akron
Youngstown
Wooster
Canton
Lima
Delaware
New Philadelphia
Newark
Springfield
Columbus
Dayton
Xenia
Chillicothe
Athens
Cincinnati
Portsmouth
Ironton
Williams
Fulton
Lake
Ashtabula
Defiance
Henry
Sandusky
Geauga
Trumbull
Lorain
Paulding
Putnam
Seneca
Medina
Van Wert
Wyandot
Crawford
Ashland
Columbiana
Mercer
Auglaize
Hardin
Marion
Holmes
Carroll
Jefferson
Shelby
Knox
Darke
Champaign
Guernsey
Belmont
Muskingum
Preble
Fairfield
Noble
Monroe
Perry
Morgan
Fayette
Washington
Butler
Clinton
Hocking
Vinton
Highland
Meigs
Brown
Adams
Jackson
Lawrence

TFN Financing

- $21M investment
- Financing from Ohio State University
  - Loan for fiber ($7M)
  - Short-term financing ($2M)
- Financing from state capital budget ($8.5M)
  - Equipment
  - Last mile to 17 institutions

Equipment

- Cisco 15454 integrated solution (DWDM)
  - all of the amps, mux/demux etc. integrated
- Multi Service Transport Platform (MSTP)
  - ITU G.709 compliant
- Cisco routers (GSR 12000) and switches
- Juniper M7i routers

Implementation

- Using MPLS
- 2.5 gig backbone
- Hired 2 optical engineers
- Using Cisco Transport Manager software
- First fiber cut tested redundancy

Video and Advanced Services as a Demonstration of TFN

Video and Advanced Services as a Demonstration of TFN

- H.323 Video
  - IntraState Service
  - Commons
  - Testing programs
- Grid Video Services
- Satellite Applications
- Ubiquitous deployment is stated goal

IntraState Service

- Provide MCU service at cost to campuses
- Provide Site Certification
- Work through issues with campuses
- Developed H.323 Beacon for end users
- 24/7 NOC
- Connections up 768Kbps
- Streaming/Archiving service

I2 Commons

- Provide MCU services
- Provide Site Certification
- Work through issues with networks/GigaPOPs
- Developed H.323 Beacon for end users
- 24/7 NOC
- Connections up 768Kbps
- Streaming/Archiving service
- Testing service on various applications sponsored by I2

Business Model

- Need to assess equipment size i.e.how many users need to be supported
- Determine if it is cheaper to buy own equipment or subscribe to service
- What is the enterprise/organizational needs

Equipment Needed

- MCU
- Gatekeeper
- Video clients to support service

Equipment Needed
High End System

- MCU - $259,894
- Gatekeeper - $327,097
- Video clients to support service - $93,136
- Total $734,867

Equipment Needed
Low End System

- MCU 20 user - $58,000
- Gatekeeper - $17,912
- Video clients to support service - $93,136
- Total $169,048

Our Cost Recovery Model

- $400 per month i.e. $4,800
- 30 conferences a month
- $400 plus $50 for every site over 2 when over 30 conferences
- Basic cost per conference = $13.34

Example

- Let's eliminate video client cost as a sunk cost i.e.no matter how you do this you need those
- High end needs to run ~44,000 conferences to break even
- need 122 minimum users

Example 2

- Same basic premise
- Low end needs to run ~5,691conferences to break even
- need 16 minimum users

Conclusion

- Need to carefully assess needs
- One needs to also consider HR overhead we did not look at that
- To run a large video conferencing operation takes significant resources

H.323 Beacon

Basics of Voice and Video over IP performance measurement

- Voice and Video Traffic are inherently different from Data Traffic
  - They use Signaling protocols such as H.323, SIP, …
  - Media (Voice and Video payload) is delivered over IP using RTP packets irrespective of which codec or end-point technology (PC-based Vs Appliance-based)
- General ICMP and UDP based tools fail to totally capture the performance bottlenecks faced by actual voice and video traffic at the end-host and in the network
  - i.e.- ping, traceroute, Iperf, pathrate, … don't suffice!
  - We need to measure end-user experience of Voice and Video over IP applications also!

Measuring End-user experience…

- Two approaches to evaluating end-user experience of audiovisual quality
  - ***Subjective Measurements***
    - Involve human participants to rate audiovisual quality ***Can you hear me now?***
    - Mean Opinion Score (MOS) Ranking technique (ITU-T P.800) ***Not just "Good"!***
  - ***Objective Measurements***
    - Automated techniques to rate audiovisual quality
      - "E-Model" [ITU-T G.107]
      - Perceptual Evaluation of Speech Quality (PESQ) [ITU-T P.862]

Mean Opinion Score (MOS) Concept

Quality Scale	Score	Listening Effort Scale
Excellent	5	No effort required
Good	4	No appreciable effort required
Fair	3	Moderate effort required
Poor	2	Considerable effort required
Bad	1	No meaning understood with reasonable effort

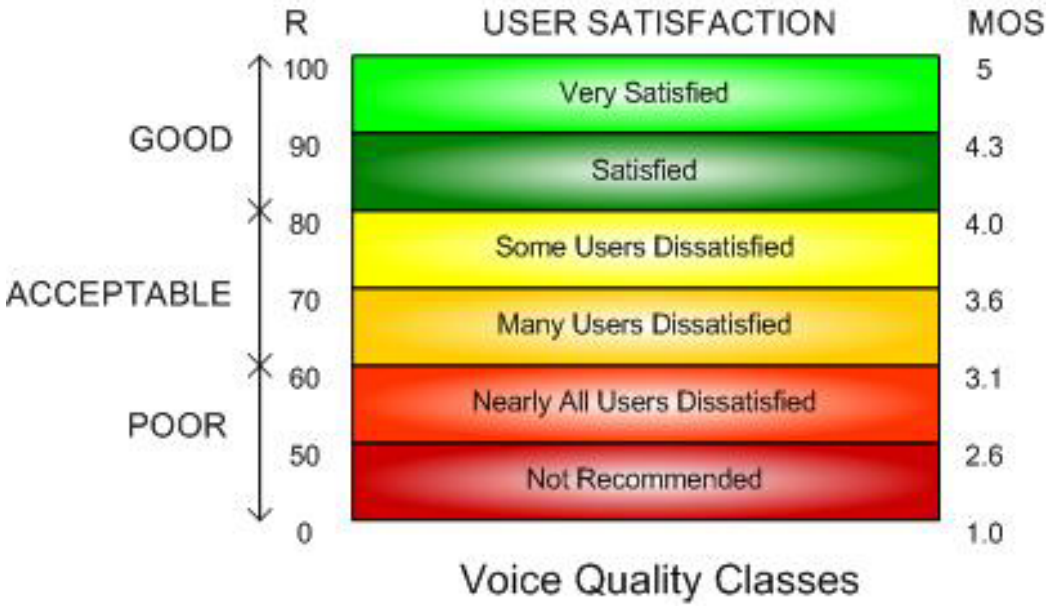

What are the other common end-to-end performance problems?

- Common problems involving endpoint devices
  - Failure of audio and video hardware, out dated or buggy end-point application software
  - Faulty connections and configurations of audio and video interfaces
  - Lack of lip-synchronization
  - Mis-configured jitter buffer sizes
  - Non inter-operable end-point application software
  - Lack of forward error correction and echo cancellation mechanisms in end-point clients
  - Lack of end-user training

Common End-to-End performance problems in VVoIP Systems (2)

- Common problems involving network devices
  - Insufficient network capacity for handling multiple high data rate videoconferencing calls
  - Excessive delay, loss, jitter, out of order packets and re-ordered packets in the network
  - Duplex mis-match problems
  - Traffic congestion at peak-usage periods of the network
  - Mis-configured 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

Common End-to-End performance problems in VVoIP Systems (3)

- Common problems involving application service devices
  - Misconfigured firewalls that block required ports
  - Non-H.323 friendly NATs
  - Misconfigured devices
    - MCUs
    - Gatekeeper
    - Gateway devices
  - Outdated or buggy application-service device application software

ActiveMON

- Generalized version of H.323 beacon

ActiveMon Architecture

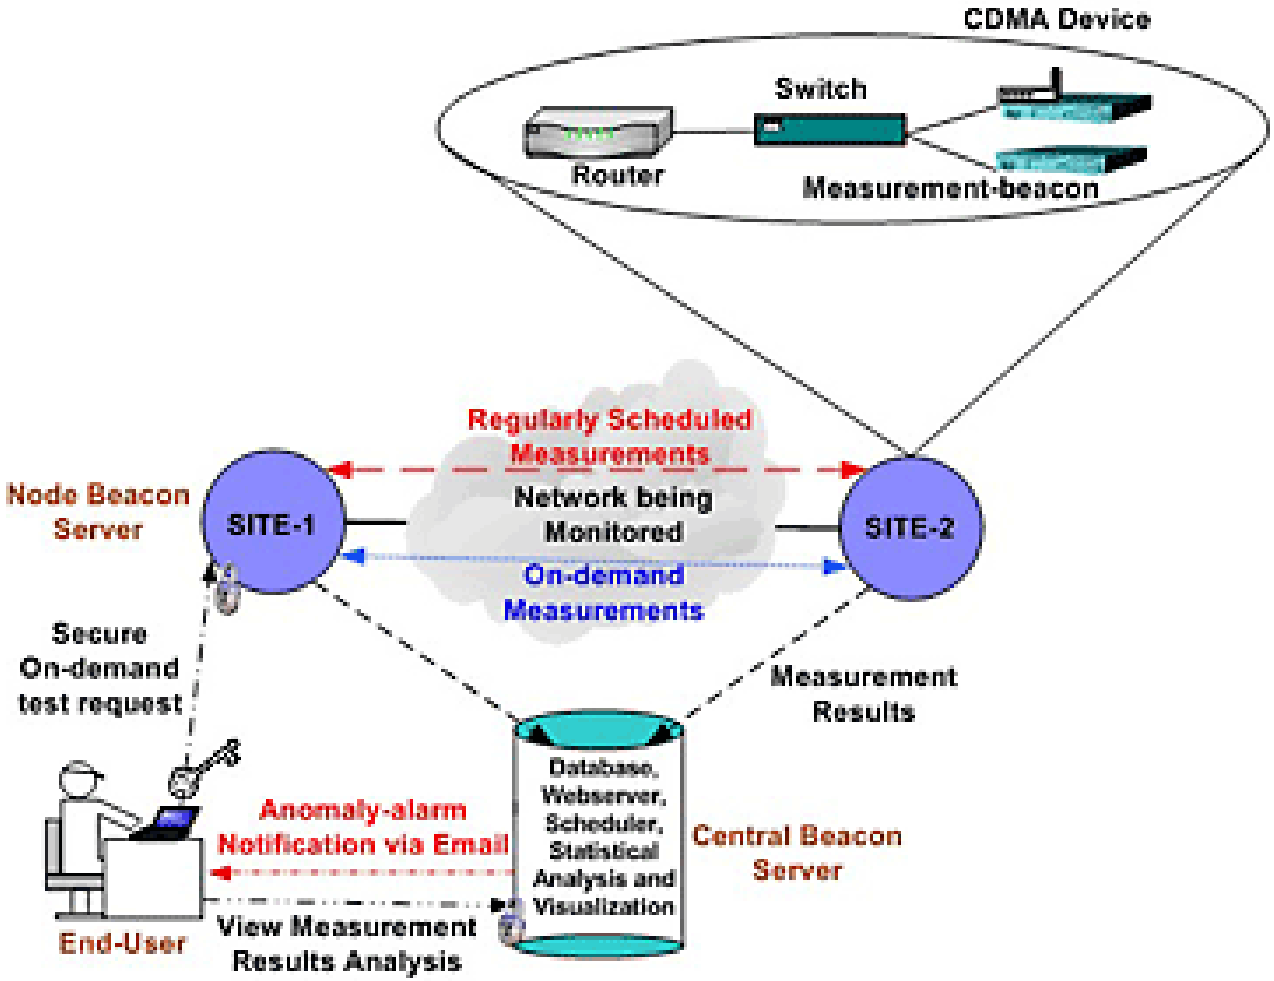

ActiveMon Framework Features

- *Data-Generator Module* for an application-specific network measurement toolkit
- *Central Data-Collector-Sanitizer Module* to centrally collect and store sensible measurements data; E.g. this module avoids collecting ‘–ve’ Delay or MOS values, etc.
- *Optimized Database Schema* to efficiently store massive amounts of measurement data with minimal redundancy; saves disk space and facilities quicker data mining

ActiveMon Measurement Toolkit

Measured Characteristics	Tool
Round-trip delay	Ping
High-precision one-way delay	OWAMP
Topology and route changes	Traceroute
Bandwidth capacity: Per-hop	Pathchar
Available bandwidth	Pathload
Bottleneck bandwidth	Pathrate
UDP transfer bandwidth, Jitter and Loss	Iperf
Performance of interactive audio/video streams (MOS)	H.323 Beacon

ActiveMon can be easily enhanced to support other tools as well…

ActiveMon Framework Features (2)

- *Scalable Scheduler Module* for handling network-wide on-going and on-demand measurements; scheduling supports regulation and prevents measurement conflicts due to resource sharing
- *Alarm Generator Module* digests, analyzes and generates alarms based on an efficient anomaly detection scheme that aims at minimum false-alarms; alarm notification via e-mail is supported
- Easily *Customizable Visualization Module* with tabular and network health Weather map interfaces; alarm-context sensitive coloring of measurements information is supported
- *Security Configurations* to avoid compromise of measurement infrastructure resources

Is ActiveMon fully developed and available as open-source?

- An alpha version of the software with several of the above features has been developed and deployed on a measurement testbed
- Based on the deployment experiences, the alpha version is being enhanced to provide better and more consistent functionality
- To obtain the alpha version of ActiveMon, please contact-

Prasad Calyam

pcalyam@oar.net

Measurements Testbed

- Goal-1:To study end-to-end network performance measurement data reported by various tools to empirically correlate network events and measurement data anomalies in a routine monitoring infrastructure

"Do measurement tools actually detect significant network events?"

- Goal-2: To analyze long-term network performance trends via statistical analysis of active and passive measurement data collected at strategic points on an ongoing basis

"What can be understood from long-term network measurements?"

- Goal-3: To use findings obtained from fulfilling the above Goals 1 and 2, to comprehensively compare performance at campus, regional and national network backbone levels and hence to quantify end-to-end network performance stability in typical hierarchical network backbones

"How does it matter where I measure the network?"

Testbed spanning Hierarchical Network Backbone Levels – Campus, Regional, National

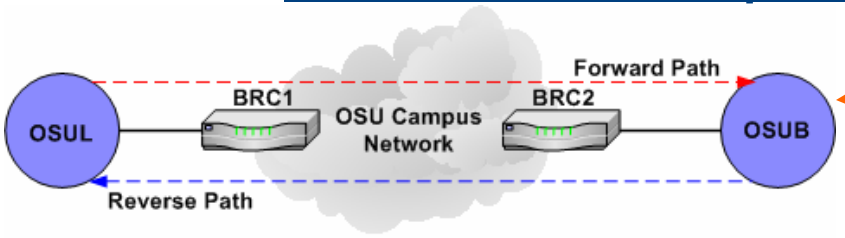

OSUL: Ohio State University Lab Router Measurement Point
OSUB: Ohio State University Border Router Measurement Point

Campus - Level Path

Only OSU Campus Backbone Routers were present along the path

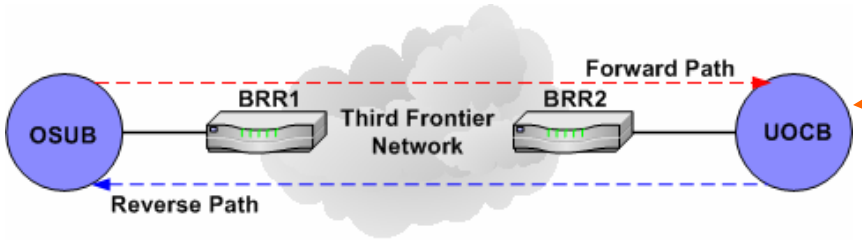

OSUB: Ohio State University Border Router Measurement Point
UOCB: University of Cincinnati Border Router Measurement Point

Regional - Level Path

Only OARnet Backbone Routers were present along the path

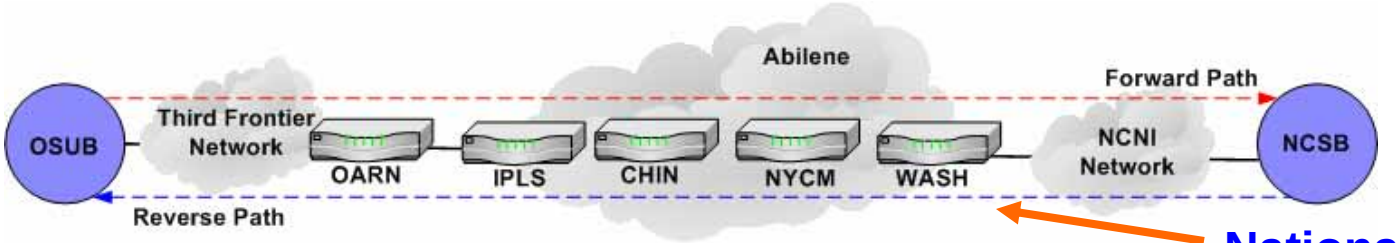

OSUB: Ohio State University Border Router Measurement Point
NCSB: North Carolina State University Border Router Measurement Point

National - Level Path

Only OARnet Backbone Routers, Abilene Routers, NCNI Routers were present along the path

Case Study - I

(July 2004 – December 2004 Measurements Data)

- **Delay Variations**
  - We found that combined one-way delays (A→B+B→A) along a path with ends A and B are comparable to round trip delays (A↔B) in all the three paths
  - Significant anomalies due to route changes (each time!)
  - Short-lived dips and peaks due to miscellaneous temporal network dynamics; Magnitudes based on hop-count

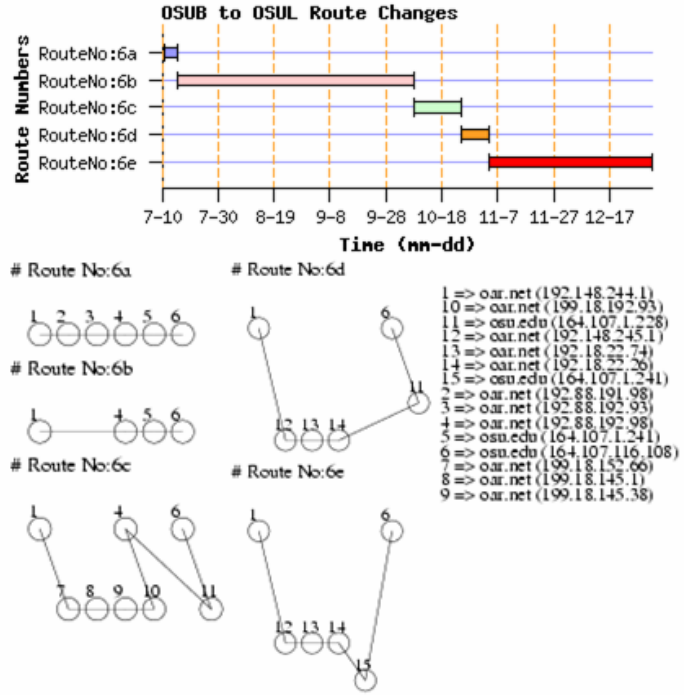

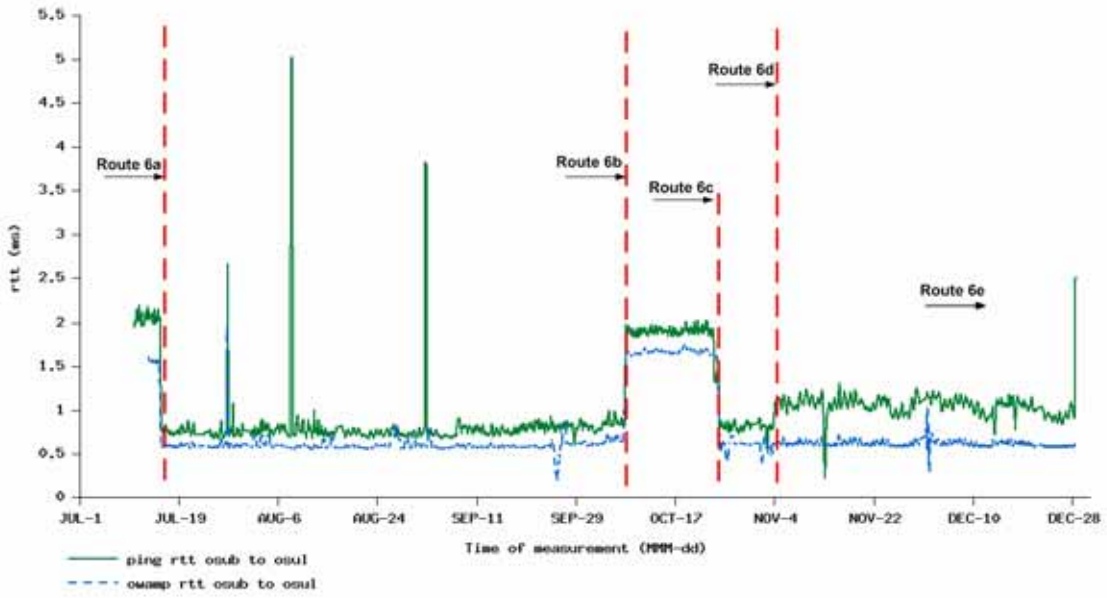

Case Study - II

(July 2004 – December 2004 Measurements Data)

- **Bandwidth Variations**
  - Router mis-configuration anomaly with three distinct trends
  - Regional path was the least congested and most provisioned path
  - National path traffic spanning multiple-ISPs experiences most congestion events

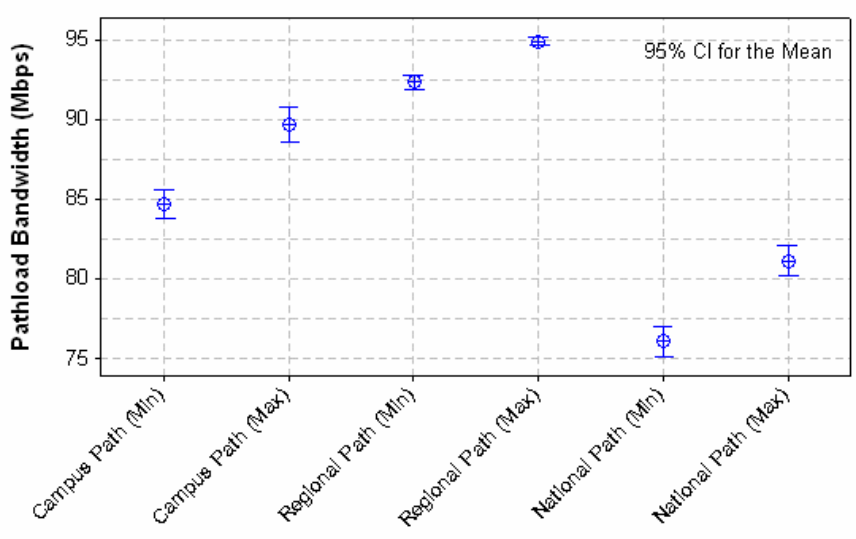

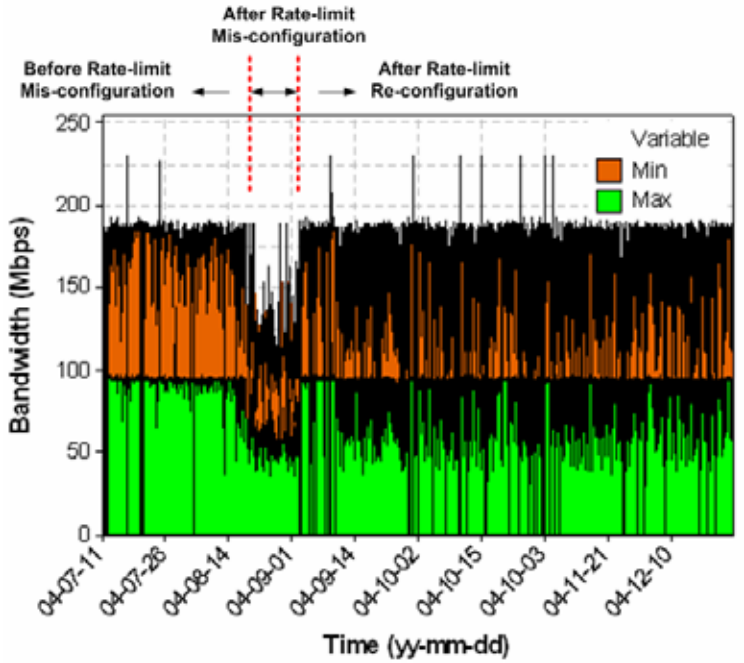

ActiveMon Related Publications (2005)

- Prasad Calyam, Dima Krymskiy, Mukundan Sridharan, Paul Schopis, "**TBI: End-to-End Network Performance Measurement Testbed for Empirical-bottleneck Detection**", ***IEEE TRIDENTCOM***, 2005.
- Prasad Calyam, Chang-Gun Lee, Phani Kumar Arava, Dima Krymskiy, David Lee, "**OnTimeMeasure: A Scalable Framework for scheduling active measurements**", ***IEEE E2EMON***, 2005.
- Prasad Calyam, Dima Krymskiy, Mukundan Sridharan, Paul Schopis, "**Active and Passive Measurements on Campus-level, Regional-level and National-level Network Backbone Paths**", ***IEEE ICCCN***, 2005.
- Prasad Calyam, Chang-Gun Lee, Phani Kumar Arava, Dima Krymskiy, "**Enhanced EDF Scheduling Algorithms for Orchestrating Network-wide Active Measurements**", ***IEEE RTSS***, 2005.

Above papers are available at - http://www.osc.edu/research/networking/publications.shtml

Grid Video Services

Grid Video Services

- Doesn't use MCU
- Achieves multipoint functionality via multicast
- Multicast has its own unique set of problems
- Multicast uses Reverse Path Forwarding(RPF)
- RPF path does not necessarily follow unicast route

The Fish Problem

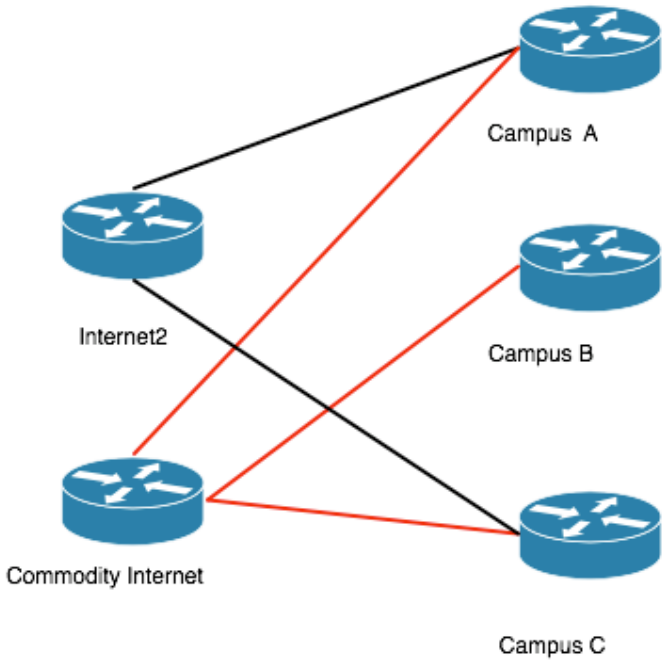

Logical Router Solution to The Fish Problem

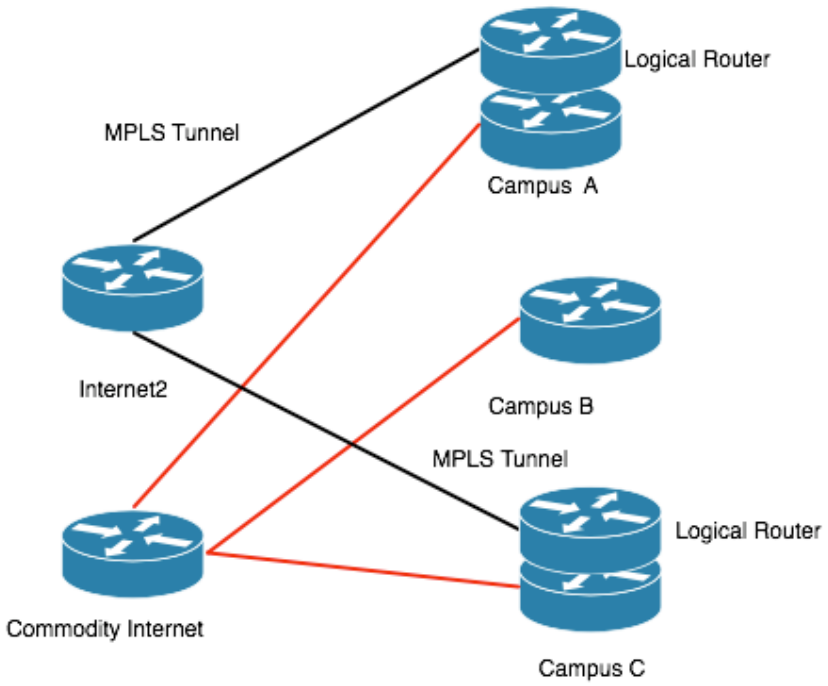

Multicast Complication in MPLS Environment

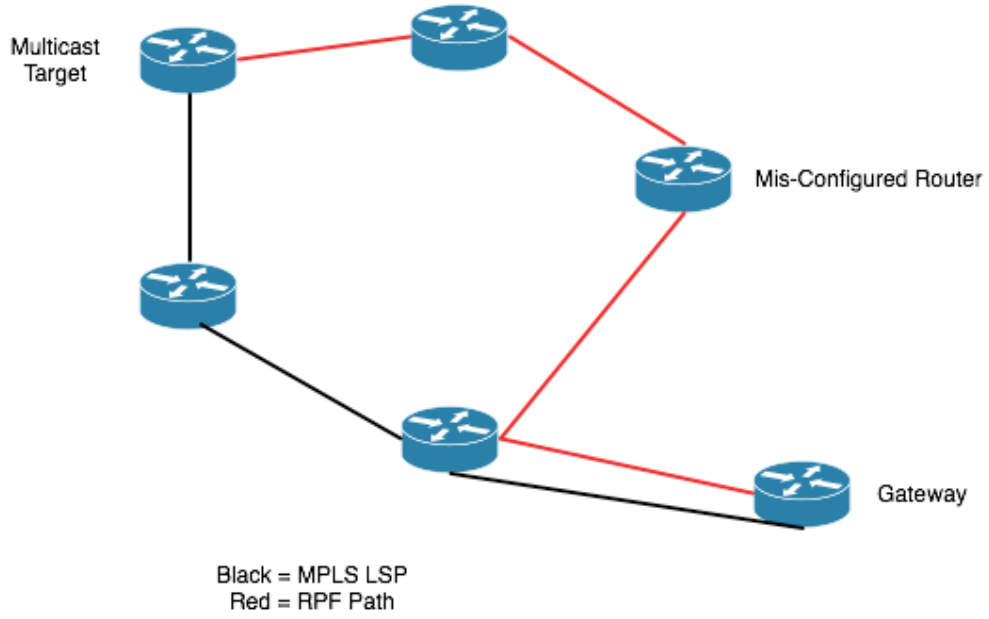

Key Points to Take Home

1. Advanced Services require additional expertise
2. Control over the infrastructure doesn't necessarily mean more control; it does give one additional information, at times overwhelmingly
3. Solutions can be slow in coming e.g. just because one understands the problem does not necessarily mean one readily has a solution

IPCOMM2006
September 25-27 • Gaylord Opryland • Nashville, TN

QUESTIONS?

Contact:
Paul Schopis
pschopis@oar.net

OSC
Innovations in computing, networking, and education

OARnet
Your Internet Solutions Partner

TFN
OHIO LIGHTS THE WAY