

Terapaths: DWMI: Datagrid Wide Area Monitoring Infrastructure

Les Cottrell, SLAC

*Presented at DoE PI Meeting BNL September
2005*

[www.slac.stanford.edu/grp/scs/net/talk05/dwmi-
sep05.ppt](http://www.slac.stanford.edu/grp/scs/net/talk05/dwmi-sep05.ppt)



Partially funded by DOE/MICS for Internet End-to-end
Performance Monitoring (IEPM)

Goals

- Develop/deploy/use a high performance network monitoring tailored to HEP needs (tiered site model):
 - Evaluate, recommend, integrate best measurement probes including for ≥ 10 Gbps & dedicated circuits
 - Develop and integrate tools for long-term forecasts
 - Develop tools to detect significant/persistent loss of network performance, AND provide alerts
 - Integrate with other infrastructures, share tools, make data available




Using Active IEPM-BW measurements

- Focus on high performance for a few hosts needing to send data to a small number of collaborator sites, e.g. HEP tiered model
- Makes regular measurements with tools, now supports
 - Ping (RTT, connectivity), traceroute
 - pathchirp, ABwE, pathload (packet pair dispersion)
 - iperf (single & multi-stream), thrulay,
 - Bbftp, bbcp (file transfer applications)
 - Looking at GridFTP but complex requiring renewing certificates
- Lots of analysis and visualization
- Running at major HEP sites: CERN, SLAC, FNAL, BNL, Caltech to about 40 remote sites
 - http://www.slac.stanford.edu/comp/net/iepm-bw.slac.stanford.edu/slac_wan_bw_tests.html

- Improved management: easier install/updates, more robust, less manual attention
- Visualization (new plots, MonALISA)
- **Passive needs & progress**
 - Packet pair problems at 10Gbits/s, timing in host and NIC offloading
 - Traffic required for throughput (e.g. > 5GBytes)
 - Evaluating effectiveness of using passive (Netflow)
 - No passwords/keys/certs, no reservations, no extra traffic, real applications, real partners...
 - ~30K large (>1MB) flows/day at SLAC border with ~ 70 remote sites
 - 90% sites have no seasonal variation so only need typical value
 - In a month 15 sites have enough flows to use seasonal methods
 - Validated that results agree with active, flow aggregation easy

But



- Apps use dynamic ports, need to use indicators to ID interesting apps
- Throughputs often depend on non-network factors:
 - Host interface speeds (DSL, 10Mbps Enet, wireless)
 - Configurations (window sizes, hosts)
 - Applications (disk/file vs mem-to-mem)
- Looking at distributions by site, often multi-modal
 - Provide medians, IQRs and max etc.



Forecasting

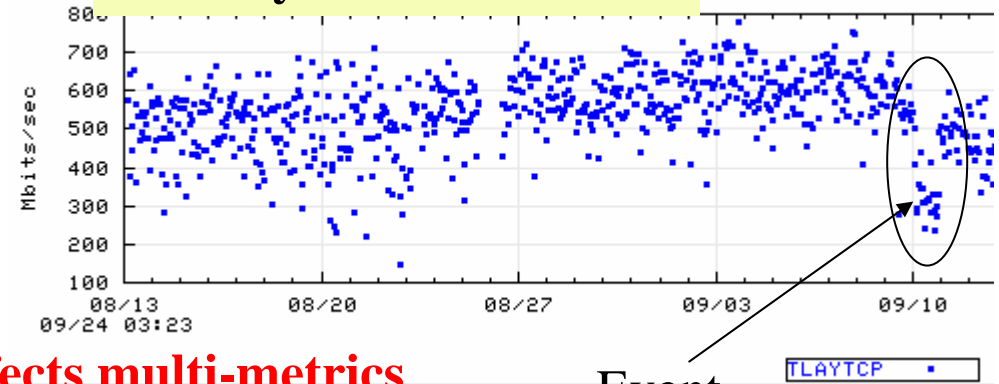
- Over-provisioned paths should have pretty flat time series
- But seasonal trends (diurnal, weekly need to be accounted for) on about 10% of our paths
- Use Holt-Winters triple exponential weighted moving averages

- Short/local term smoothing
- Long term linear trends
- Seasonal smoothing



Event detection

Thrulay SLAC to Caltech

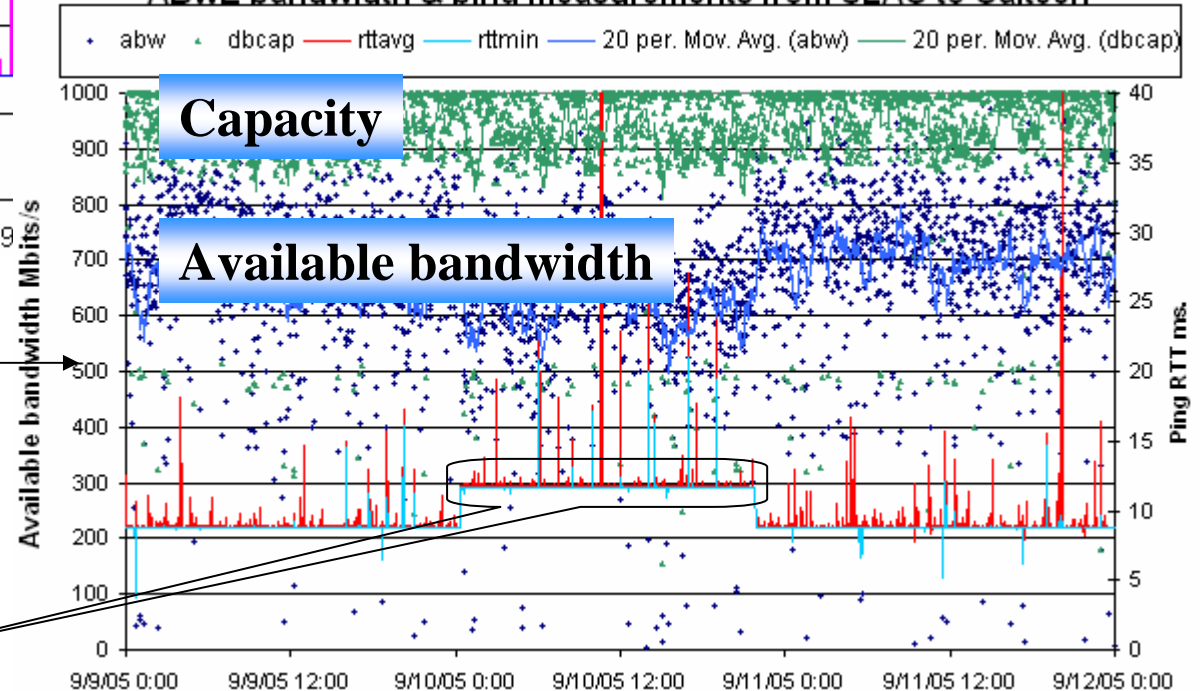


Affects multi-metrics

Event

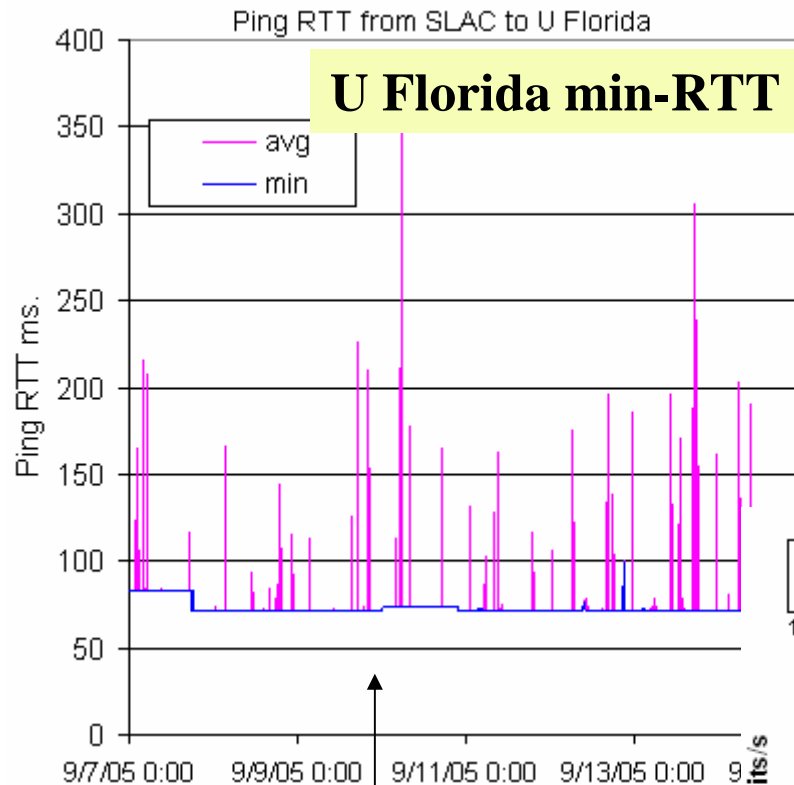
Packet pair & ping RTT

ABwE bandwidth & ping measurements from SLAC to Caltech



Affects multi-paths

Change in min-RTT



Alerts, e.g.

- Often not simple, simple RTT steps often fail:
 - <5% route changes cause noticeable thruput changes
 - ~40% thruput changes NOT associated with route change
- Use multiple metrics
 - User cares about throughput **SO** need iperf/thrulay &/or a file transfer app, **BUT** heavy net impact
 - Packet pair available bandwidth, lightweight but noisy, needs timing (hard at > 1Gbits/s and TCP Offload in NICs)
 - Min ping RTT & route changes may have no effect on throughput
- Look at multiple routes
- Fixed thresholds poor (need manual setting), need automation
- Some routes have seasonal effects



Collaborations



- HEP sites: BNL, Caltech, CERN, FNAL, SLAC, NIIT
- ESnet/OSCARS – Chin Guok
- BNL/QoS- Dantong Yu
- Development – Maxim Grigoriev/FNAL, NIIT/Pakistan
- Integrate our traceroute analysis/visualization into AMP (NLNR) – Tony McGregor
- Integrate IEPM measurements into MonALISA – Iosif Legrand/Caltech/CERN

More Information

- Case studies of performance events
 - www.slac.stanford.edu/grp/scs/net/case/html/
- IEPM-BW site
 - www-iepm.slac.stanford.edu/
 - www.slac.stanford.edu/comp/net/iepm-bw.slac.stanford.edu/slac_wan_bw_tests.html
- OSCARS measurements
 - <http://www-iepm.slac.stanford.edu/dwmi/oscars/>
- Forecasting and event detection
 - www.acm.org/sigs/sigcomm/sigcomm2004/workshop_papers/nts26-logg1.pdf
- Traceroute visualization
 - www.slac.stanford.edu/cgi-wrap/pubpage?slac-pub-10341
- <http://monalisa.cacr.caltech.edu/>
 - Clients=>MonALISA Client=>Start MonALISA GUI => Groups => Test
=> Click on IEPM-SLAC



Extra Slides

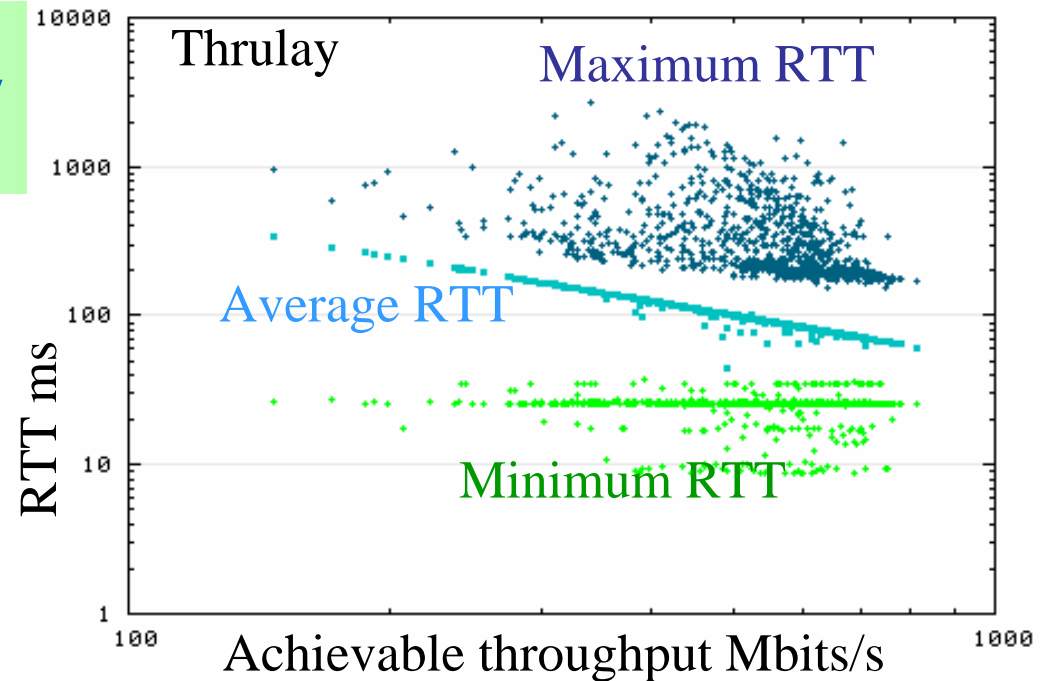


Achievable Throughput


- Use TCP or UDP to send as much data as can memory to memory from source to destination
- Tools: iperf (bwctl/I2), netperf, thrulay (from Stas Shalunov/I2), udpmon ...
- Pseudo file copy: Bbcp and GridFTP also have memory to memory mode

Iperf vs thrulay

- Iperf has multi streams
- Thrulay more manageable & gives RTT
- They agree well
- Throughput $\sim 1/\text{avg}(\text{RTT})$



TCP achievable throughput using iperf and thrulay from SLAC to Caltech



BUT...

- At 10Gbits/s on transatlantic path Slow start takes over 6 seconds
 - To get 90% of measurement in congestion avoidance need to measure for 1 minute (5.25 GBytes at 7Gbits/s (today's typical performance))
- Needs scheduling to scale, even then ...
- It's not disk-to-disk or application-to application
 - So use bbcp, bbftp, or GridFTP

- For testbeds such as UltraLight, UltraScienceNet etc. have to reserve the path
 - So the measurement infrastructure needs to add capability to reserve the path (so need API to reservation application)
 - OSCARS from ESnet developing a web services interface (<http://www.es.net/oscars/>):
 - For lightweight have a “persistent” capability
 - For more intrusive, must reserve just before make measurement

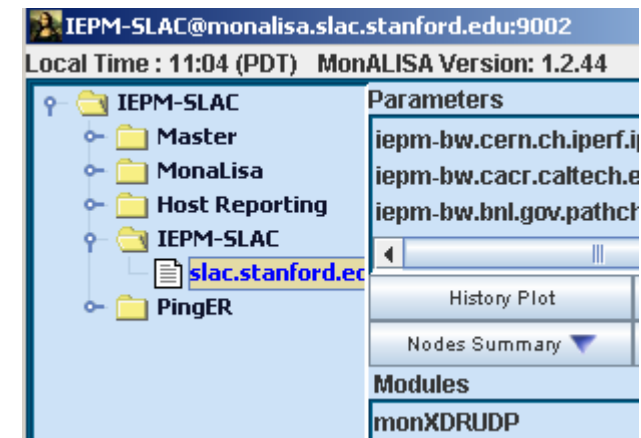

Visualization & Forecasting



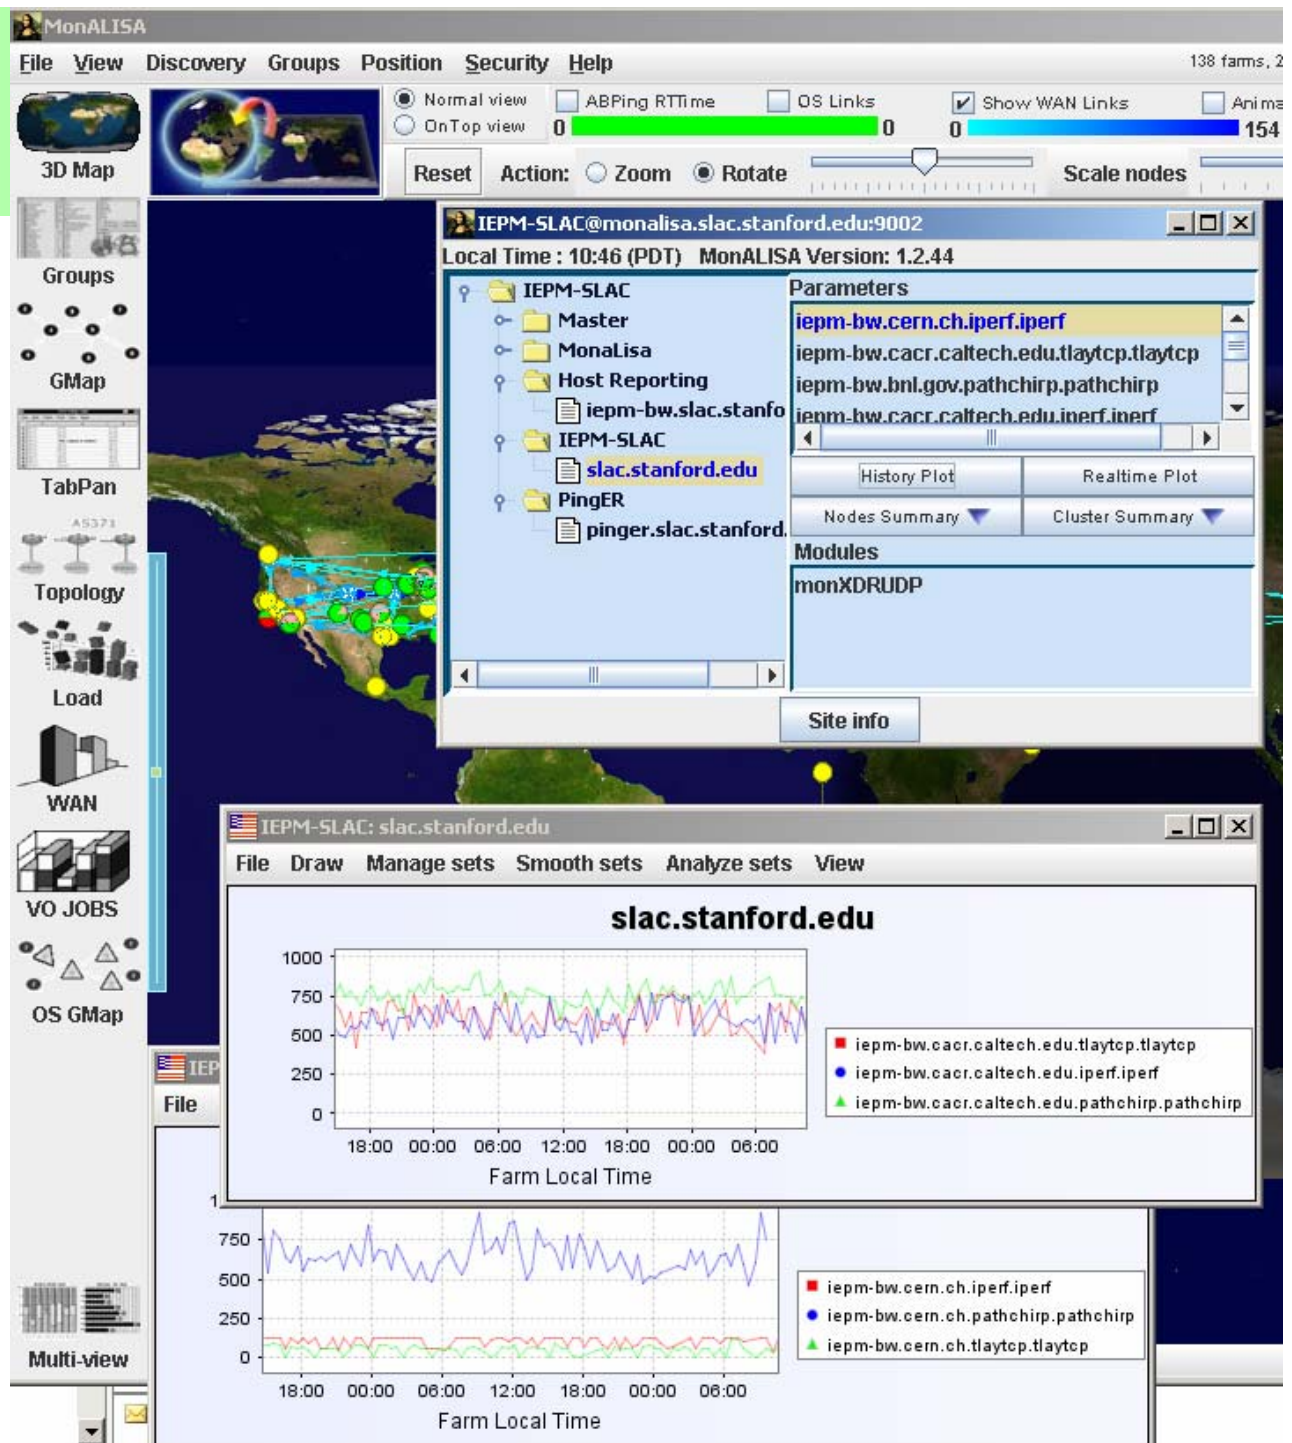
Visualization



- MonALISA (monalisa.cacr.caltech.edu/)
 - Caltech tool for drill down & visualization
 - Access to recent (last 30 days) data
 - For IEPM-BW, PingER and monitor host specific parameters
 - Adding web service access to ML SLAC data
- <http://monalisa.cacr.caltech.edu/>
 - Clients=>MonALISA Client=>Start MonALISA GUI => Groups => Test => Click on IEPM-SLAC



ML example





Changes in network topology (BGP) can result in dramatic changes in performance

Traceroute Analysis for 10/09/2003

[Yesterday's Summary](#) | [Reverse Traceroute Summary](#) | [Directory of Historical Traceroutes](#)

Hour (PST) →

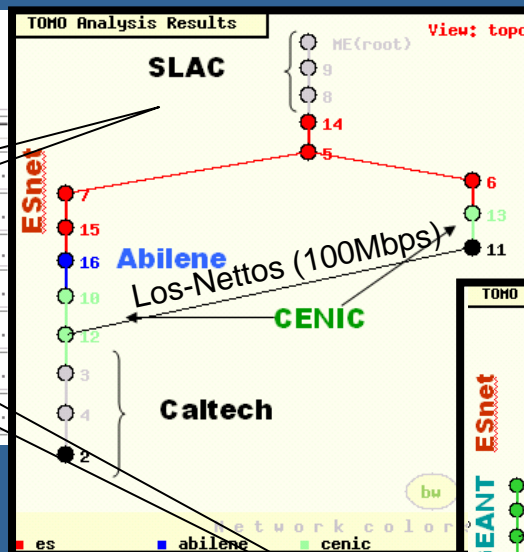
Remote host

NODE \ Hour =>	00	01	02	14	15	16	17
<input checked="" type="checkbox"/> node1.cacr.caltech.edu* R Sum Log*	105	110	...	105	...
<input type="checkbox"/> node1.cesnet.cz* R Sum Log*	35	56	35
<input type="checkbox"/> node1.clrc.ac.uk* R Sum Log*	67	71	67
<input type="checkbox"/> node1.dl.ac.uk* R Sum Log*	97	102	97
<input type="checkbox"/> node1.ece.rice.edu* R Sum Log*	104
<input type="checkbox"/> node1.fnal.gov* R Sum Log*	8
<input type="checkbox"/> node1.in2p3.fr* R Sum Log*	29	77	100


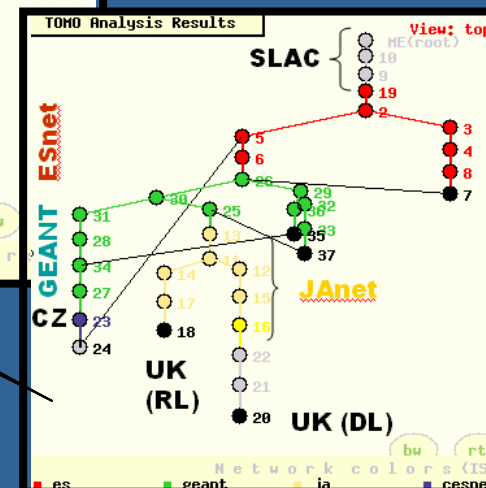
Snapshot of traceroute summary table

Notes:

1. Caltech misrouted via Los-Nettos 100Mbps commercial net 14:00-17:00
2. ESnet/GEANT working on routes from 2:00 to 14:00
3. A previous occurrence went un-noticed for 2 months
4. Next step is to auto detect and notify



Samples of traceroute trees generated from the table



ABwE measurement one/minute for 24 hours Thurs Oct 9 9:00am to Fri Oct 10 9:01am

Alerting

- Have false positives down to reasonable level, so sending alerts
- Experimental
- Typically few per week.
- Currently by email to network admins
 - Adding pointers to extra information to assist admin in further diagnosing the problem, including:
 - Traceroutes, monitoring host parms, time series for RTT, pathchirp, thrulay etc.
 - Plan to add on-demand measurements (excited about perfSONAR)

Integration

- Integrate IEPM-BW and PingER measurements with MonALISA to provide additional access
- Working to make traceanal a callable module
 - Integrating with AMP
- When comfortable with forecasting, event detection will generalize


Passive - Netflow

Netflow et. al.

- Switch identifies flow by src/dst ports, protocol
- Cuts record for each flow:
 - src, dst, ports, protocol, TOS, start, end time
- Collect records and analyze
- Can be a lot of data to collect each day, needs lot cpu
 - Hundreds of MBytes to GBytes
- No intrusive traffic, real: traffic, collaborators, applications
- No accounts/pwds/certs/keys
- No reservations etc
- Characterize traffic: top talkers, applications, flow lengths etc.
- Internet 2 backbone
 - <http://netflow.internet2.edu/weekly/>
- SLAC:
 - www.slac.stanford.edu/comp/net/slac-netflow/html/SLAC-netflow.html


Typical day's flows

- Very much work in progress
- Look at SLAC border
- Typical day:
 - >100KB flows
 - ~ 28K flows/day
 - ~ 75 sites with > 100KByte bulk-data flows
 - Few hundred flows > GByte



Forecasting?


- Collect records for several weeks
- Filter 40 major collaborator sites, big (> 100KBytes) flows, bulk transport apps/ports (bbcp, bbftp, iperf, thrulay, scp, ftp)
- Divide by remote site, aggregate parallel streams
- Fold data onto one week, see bands at known capacities and RTTs



Peaks at known capacities and RTTs


RTTs might suggest windows not optimized

SLAC Border Netflow long flow throughputs, Jan '05





How many sites have enough flows?

- In May '05 found 15 sites at SLAC border with > 1440 (1/30 mins) flows
 - Enough for time series forecasting for seasonal effects
- Three sites (Caltech, BNL, CERN) were actively monitored
- Rest were “free”
- Only 10% sites have big seasonal effects in active measurement
- Remainder need fewer flows
- So promising




Compare active with passive

Netflow passive forecast vs. ABwE available bandwidth for May 2005
from SLAC to Padova



- Predict flow throughputs from Netflow data for SLAC to Padova for May '05
- Compare with E2E active ABwE measurements

ABwE vs. Netflow forecasts, SLAC to Padova, May '05




Netflow limitations

- Use of dynamic ports.
 - GridFTP, bbcp, bbftp can use fixed ports
 - P2P often uses dynamic ports
 - Discriminate type of flow based on headers (not relying on ports)
 - Types: bulk data, interactive ...
 - Discriminators: inter-arrival time, length of flow, packet length, volume of flow
 - Use machine learning/neural nets to cluster flows
 - E.g. <http://www.pam2004.org/papers/166.pdf>
- Aggregation of parallel flows (not difficult)
- SCAMPI/FFPF/MAPI allows more flexible flow definition
 - See www.ist-scampi.org/
- Use application logs (OK if small number)

More challenges

- Throughputs often depend on non-network factors:
 - Host interface speeds (DSL, 10Mbps Enet, wireless)
 - Configurations (window sizes, hosts)
 - Applications (disk/file vs mem-to-mem)
- Looking at distributions by site, often multi-modal
- Predictions may have large standard deviations
- How much to report to application

Conclusions

- Traceroute dead for dedicated paths
- Some things continue to work
 - Ping, owamp
 - Iperf, thrulay, bbftp ... but
- Packet pair dispersion needs work, its time may be over 
- Passive looks promising with Netflow
- SNMP needs AS to make accessible
- Capture expensive
 - ~\$100K (*Joerg Micheel*) for OC192Mon



More information

- Comparisons of Active Infrastructures:
 - www.slac.stanford.edu/grp/scs/net/proposals/infra-mon.html
- Some active public measurement infrastructures:
 - www-iepm.slac.stanford.edu/
 - e2epi.internet2.edu/owamp/
 - amp.nlanr.net/
 - www-iepm.slac.stanford.edu/pinger/
- Capture at 10Gbits/s
 - www.endace.com (DAG), www.pam2005.org/PDF/34310233.pdf
 - www.ist-scampi.org/ (also MAPI, FFPP), www.ist-lobster.org
- Monitoring tools
 - www.slac.stanford.edu/xorg/nmtf/nmtf-tools.html
 - www.caida.org/tools/
 - Google for iperf, thrulay, bwctl, pathload, pathchirp



Extra Slides Follow



Visualizing traceroutes

STANFORD LINEAR ACCELERATOR CENTER



- One compact page per day
- One row per host, one column per hour
- One character per traceroute to indicate pathology or change (usually period(.) = no change)
- Identify unique routes with a number
 - Be able to inspect the route associated with a route number

Provide for analysis of long term route evolutions

[Yesterday's Summary](#) | [Reverse Traceroute Summary](#) | [Directory of Historical Traceroutes](#)

Checking a box for a node(s) and an hour(s) and pressing SUBMIT will provide topology m


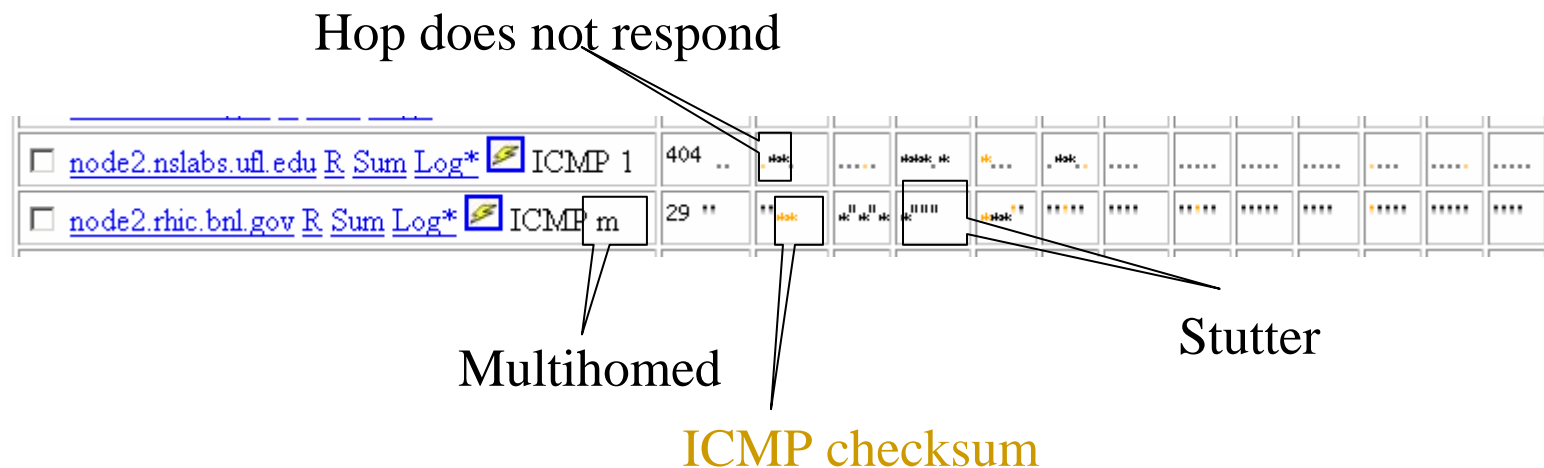

SUBMIT Topology request RESET FIELDS

NODE \ Hour (Pacific Time)=>	<input type="checkbox"/> 00	<input type="checkbox"/> 01	<input type="checkbox"/> 02	<input type="checkbox"/> 03	<input type="checkbox"/> 04
<input type="checkbox"/> node1.cacr.caltech.edu* R Sum Log*	202
<input type="checkbox"/> node1.cesnet.cz* R Sum Log*	35 ...	68	35
<input type="checkbox"/> node1.clrc.ac.uk* R Sum Log*	91 ...	112	91
<input type="checkbox"/> node1.dl.ac.uk* R Sum Log*	97 ...	155	97
<input type="checkbox"/> node1.ece.rice.edu* R Sum Log*	241
<input type="checkbox"/> node1.fnal.gov* R Sum Log*	8 ...	48	8
<input type="checkbox"/> node1.in2p3.fr* R Sum Log*	29 ...	131	30

Route # at start of day, gives idea of route stability

Multiple route changes (due to GEANT), later restored to original route

Period (.) means no change











```
tracert to CCSVSN04.IN2P3.FR (134.158.104.199), 30 hops max, 38 byte packets
 1 rtr-gsr-test (134.79.243.1) 0.102 ms
 ...
13 in2p3-lyon.cssi.renater.fr (193.51.181.6) 154.063 ms !X
```

#date	time	numnhops	epoch	rtno	route
07/08/2004	00:10:46	13	1089270646	116	(134.79.243.1), (134.79.243.1)
07/08/2004	00:25:41	14	1089271541	115	(134.79.243.1), (134.79.243.1)
07/08/2004	00:40:25	15	1089272425	114	(134.79.243.1), (134.79.243.1)
07/08/2004	00:55:24	13	1089273324	116	(134.79.243.1), (134.79.243.1)

#rt#	firstseen	lastseen	route
0	1086844945	1089705757	...,192.68.191.83,137.164.23.41,137.164.22.37,...,131.215.xxx.xxx
1	1087467754	1089702792	...,192.68.191.83,171.64.1.132,137,...,131.215.xxx.xxx
2	1087472550	1087473162	...,192.68.191.83,137.164.23.41,137.164.22.37,...,131.215.xxx.xxx
3	1087529551	1087954977	...,192.68.191.83,137.164.23.41,137.164.22.37,...,131.215.xxx.xxx
4	1087875771	1087955566	...,192.68.191.83,137.164.23.41,137.164.22.37,...,(n/a),131.215.xxx.xxx
5	1087957378	1087957378	...,192.68.191.83,137.164.23.41,137.164.22.37,...,131.215.xxx.xxx
6	1088221368	1088221368	...,192.68.191.146,134.55.209.1,134.55.209.6,...,131.215.xxx.xxx
7	1089217384	1089615761	...,192.68.191.83,137.164.23.41,(n/a),...,131.215.xxx.xxx
8	1089294790	1089432163	...,192.68.191.83,137.164.23.41,137.164.22.37,(n/a),...,131.215.xxx.xxx

[Today's Summary](#) | [Previous day's Summary](#) | [Reverse Traceroute Summary](#) | [Directory of Historical Traceroutes](#) | [Help](#)

	Parent Directory	14-Apr-2004 09:42
	2002 09/	30-Apr-2004 16:18
	2002 10/	30-Apr-2004 16:23
	2002 11/	30-Apr-2004 16:28
	2002 12/	30-Apr-2004 16:32
	2003 01/	30-Apr-2004 16:38

Character encoding of routes

- A '.' indicates that the traceroute was exactly the same as the previous one.
- A '!' indicates that the traceroute was exactly the same as the previous one, but that the datapoint is from the bw-tests regular run and not the more frequent times an hour runs.
- A '!' indicates that the traceroute was exactly the same as the previous one, but an ! annotation was found in the traceroute.
- A '|' indicates that the last hop was not reachable (i.e. the traceroute terminated after 30 hops, possibly the end host is behind a firewall).
- A red '|' indicates that the unreachable last hop, was also not pingable (probably host was unreachable).

AS' information

[Today's Summary](#) | [Previous day's Summary](#) | [Directory of Historical Traceroutes](#) | [Help](#)

☐ SUBMIT Topology request
 ☒ SUBMIT Traceroute/ASN request

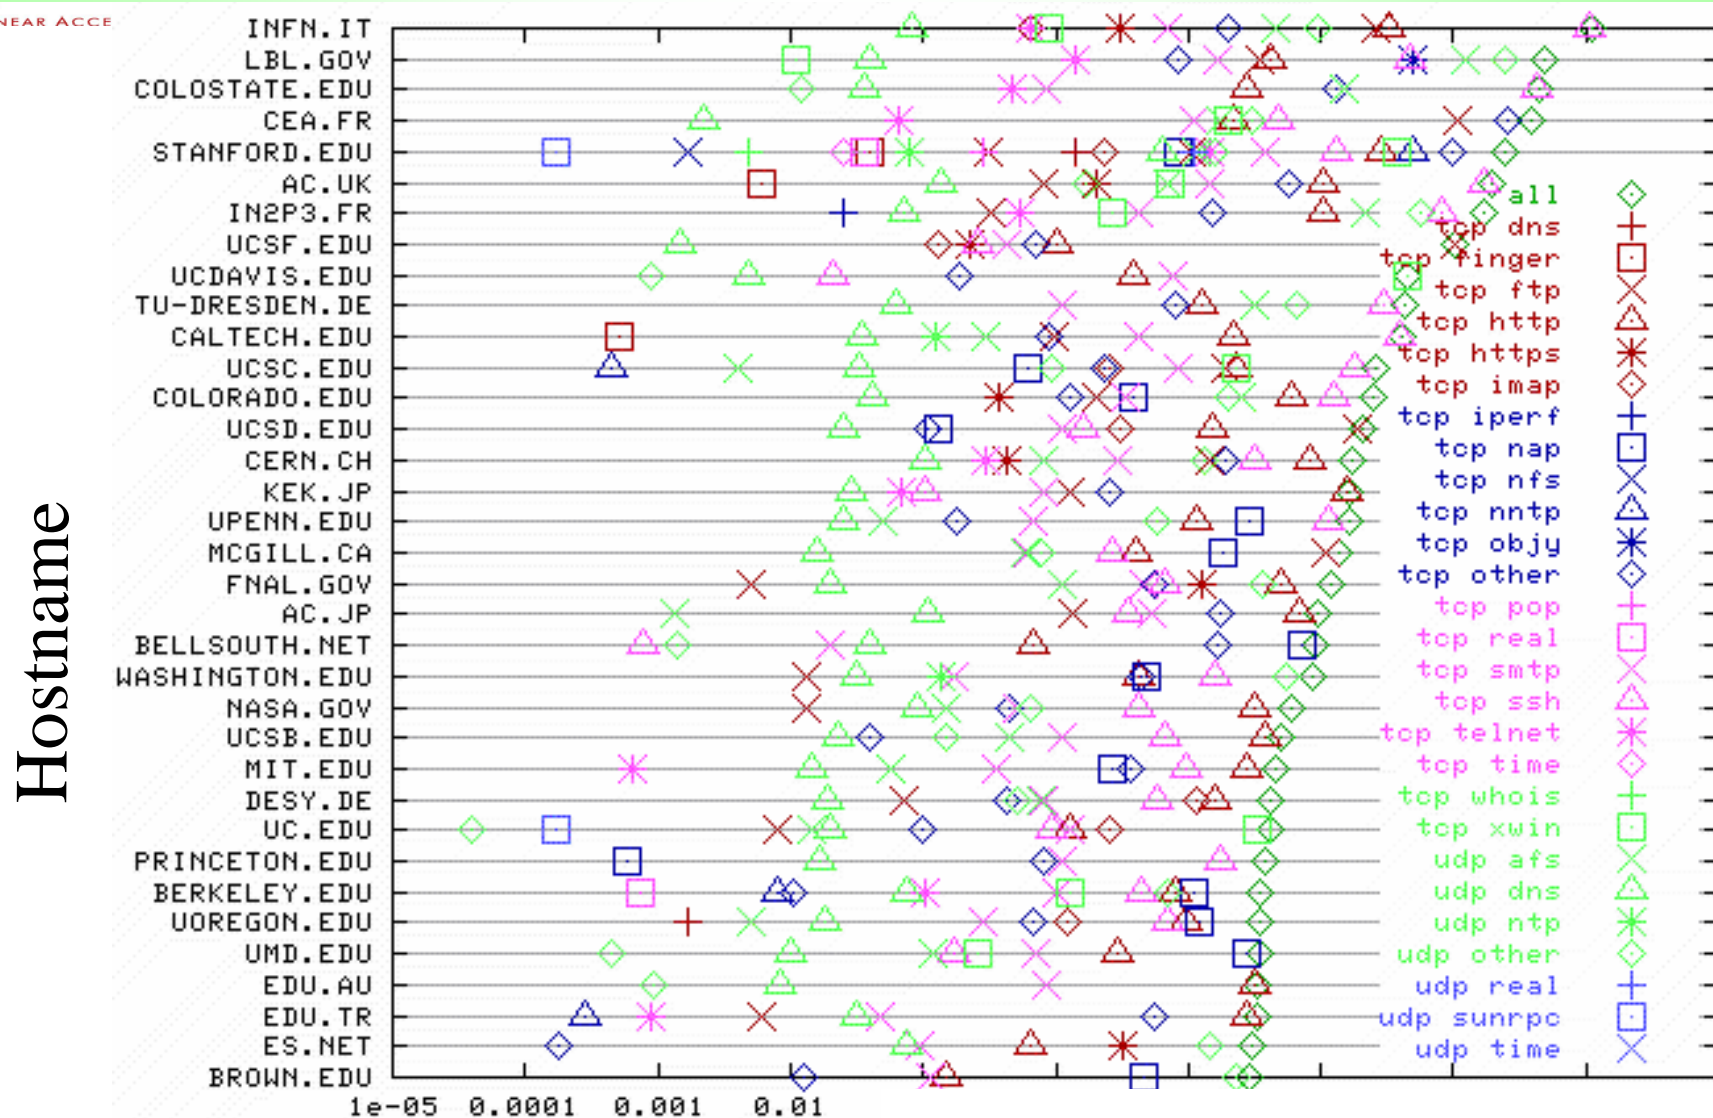
NODE \ Hour (Pacific Time)=>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
	00	01	02	03	04	05	06	07
<input checked="" type="checkbox"/> node1.binp.nsk.su R Sum Log UDP 1	0
<input type="checkbox"/> node1.cacr.caltech.edu R Sum Log UDP 1	0

traceroute to rainbow.inp.nsk.su (193.124.167.29), 30 hops max, 38 byte packets AS5402: BINP

```


1 rtr-gsr-test (134.79.243.1) 0.134 ms AS3671: SU-SLAC
2 rtr-dmz1-ger (134.79.135.15) 0.242 ms AS3671: SU-SLAC
3 slac-rt4.es.net (192.68.191.146) 0.339 ms SLAC-1: Stanford
4 snv-pos-slac.es.net (134.55.209.1) 0.933 ms AS293: Energy
5 chicr1-oc192-snvcr1.es.net (134.55.209.54) 48.989 ms AS293: Energy
6 aoacr1-oc192-chicr1.es.net (134.55.209.58) 69.059 ms AS293: Energy
7 aoapr1-ge0-aoacr1.es.net (134.55.209.110) 69.592 ms AS293: Energy
8 198.124.216.126 (198.124.216.126) 256.832 ms AS291: ESnet-CIDR-A
9 keksw2-ns.kek.jp (130.87.4.35) 266.092 ms AS2505: KEK
    
```

Top talkers by application/port



Flow sizes

Flow size distribution at SLAC border April 9, 2001



Real
A/V

AFS
file
server

Heavy tailed, in ~ out, UDP flows shorter than TCP, packet~bytes

75% TCP-in < 5kBytes, 75% TCP-out < 1.5kBytes (<10pkts)

UDP 80% < 600Bytes (75% < 3 pkts), ~10 * more TCP than UDP

Top UDP = AFS (>55%), Real(~25%), SNMP(~1.4%)

Passive SNMP MIBs

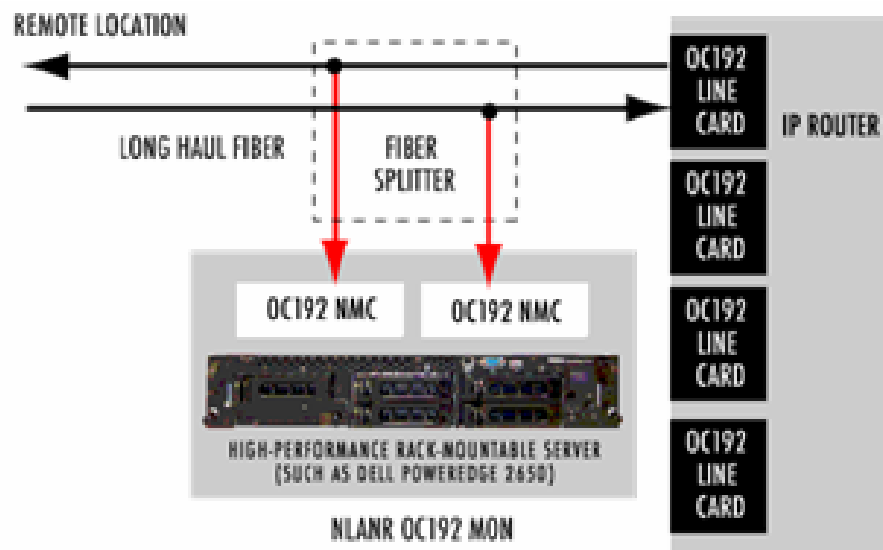
Apply forecasts to Network device utilizations to find bottlenecks

- Get measurements from Internet2/ESnet/Geant perfSONAR project
 - ISP reads MIBs saves in RRD database
 - Make RRD info available via web services
- Save as time series, forecast for each interface
- For given path and duration forecast most probable bottlenecks
- Use MPLS to apply QoS at bottlenecks (rather than for the entire path) for selected applications
- NSF proposal

Passive – Packet capture

10G Passive capture

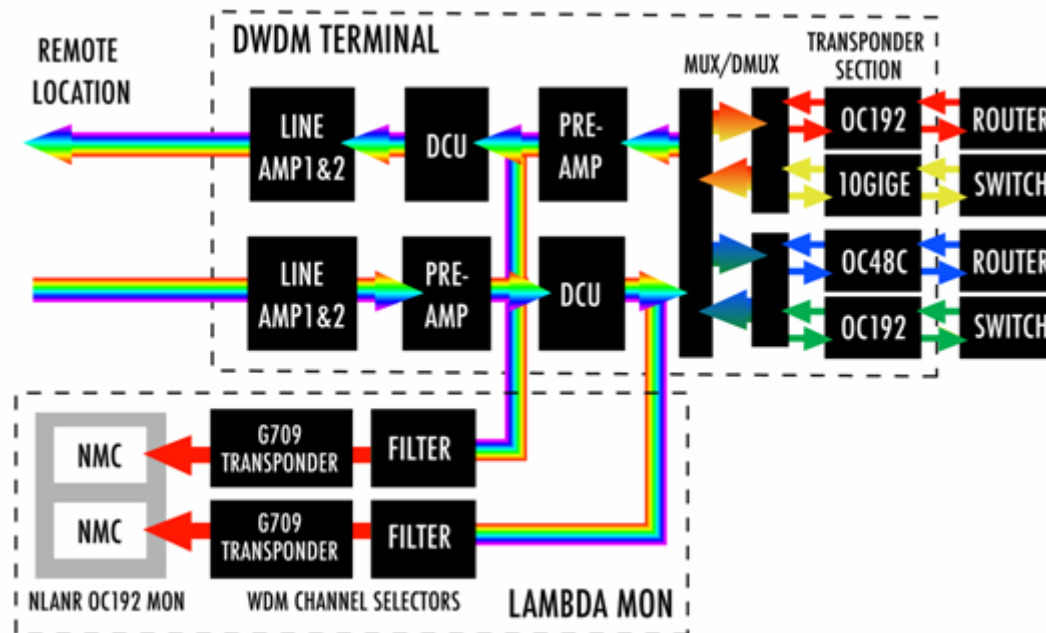
- Endace (www.endace.net): OC192 Network Measurement Cards = DAG 6 (offload vs NIC)
 - Commercial OC192Mon, non-commercial SCAMPI



- Line rate, capture up to $> \sim 1\text{Gbps}$
- Expensive, massive data capture (e.g. PB/week) tap insertion
- D.I.Y. with NICs instead of NMC DAGs
 - Need PCI-E or PCI-2DDR, powerful multi CPU host
 - Apply sampling
 - See www.uninett.no/publikasjoner/foredrag/scampi-noms2004.pdf

LambdaMon / Joerg Micheel NLANR

- Tap G709 signals in DWDM equipment
- Filter required wavelength
- Can monitor multiple λ 's sequentially



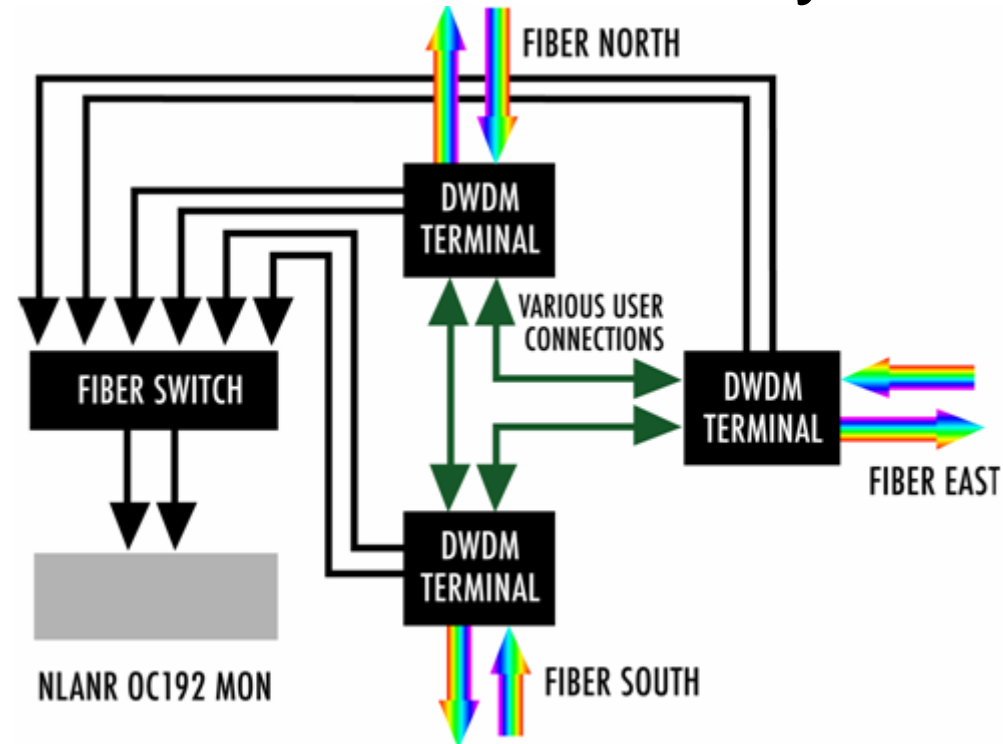
2 tunable filters



LambdaMon



- Place at PoP, add switch to monitor many fibers
- More cost effective




- Multiple G.709 transponders for 10G
- Low level signals, amplification expensive
- Even more costly, funding/loans ended ...

Ping/traceroute

- Ping still useful (*plus ca reste ...*)
 - Is path connected?
 - RTT, loss, jitter
 - Great for low performance links (e.g. Digital Divide), e.g. AMP (NLNR)/PingER (SLAC)
 - Nothing to install, but blocking
- **OWAMP/I2** similar but **One Way**
 - But needs server installed at other end and good timers
- Traceroute
 - Needs good visualization (traceanal/SLAC)
 - Little use for dedicated λ layer 1 or 2
 - However still want to know topology of paths

Packet Pair Dispersion



- Send packets with known separation
- See how separation changes due to bottleneck
- Can be low network intrusive, e.g. ABwE only 20 packets/direction, also fast < 1 sec
- From PAM paper, pathchirp more accurate than ABwE, but
 - Ten times as long (10s vs 1s)
 - More network traffic (~factor of 10)
 - Pathload factor of 10 again more
 - <http://www.pam2005.org/PDF/34310310.pdf>
- IEPM-BW now supports ABwE, Pathchirp, Pathload