

Wide Area Networking at SLAC, Feb '03

Warren Matthews, Les Cottrell, Paola Grosso

SLAC/SCS

Prepared for BaBar Data Distribution meeting, SLAC
February 26 2003

www.slac.stanford.edu/grp/scs/net/talk/babar-dd-deb03.htm

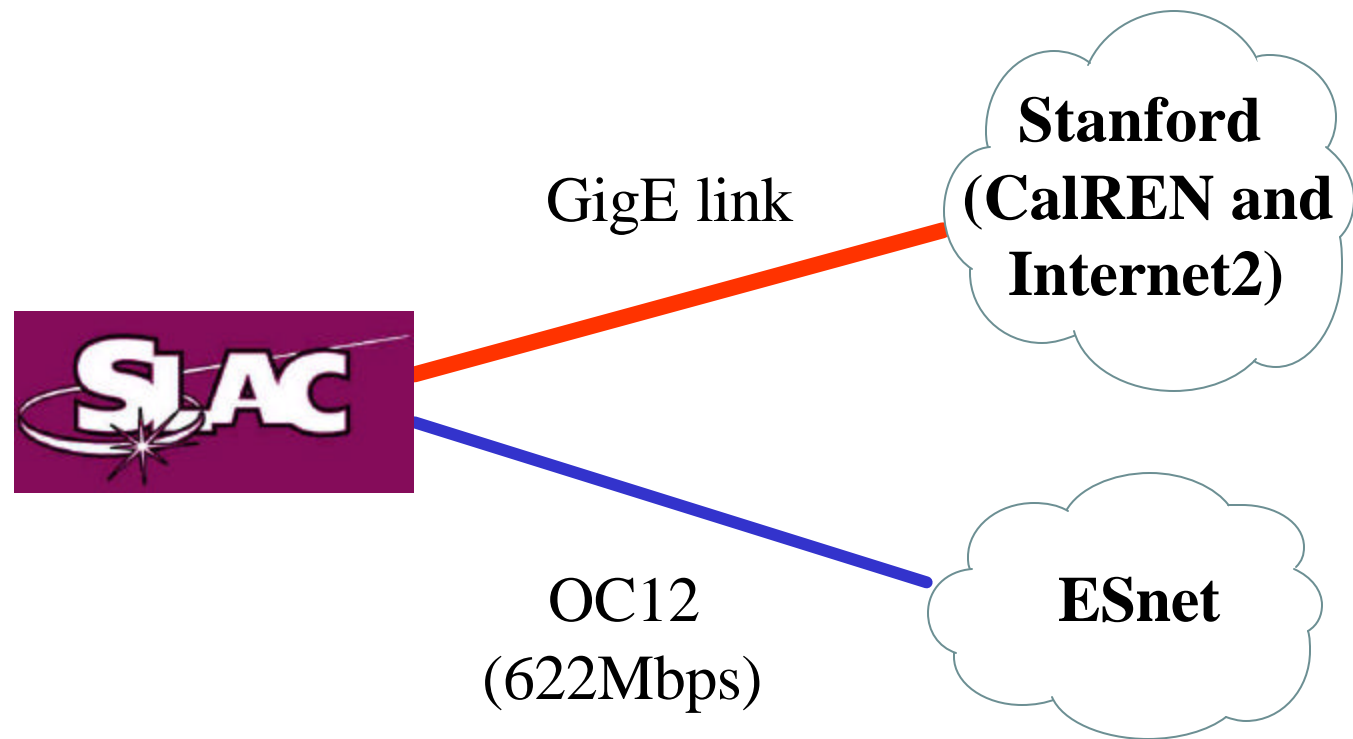


Outline

- SLAC Connections to the Internet
- Usage
- Performance



WAN Connections



These networks connect (peer) with other networks providing SLAC connectivity to the world.



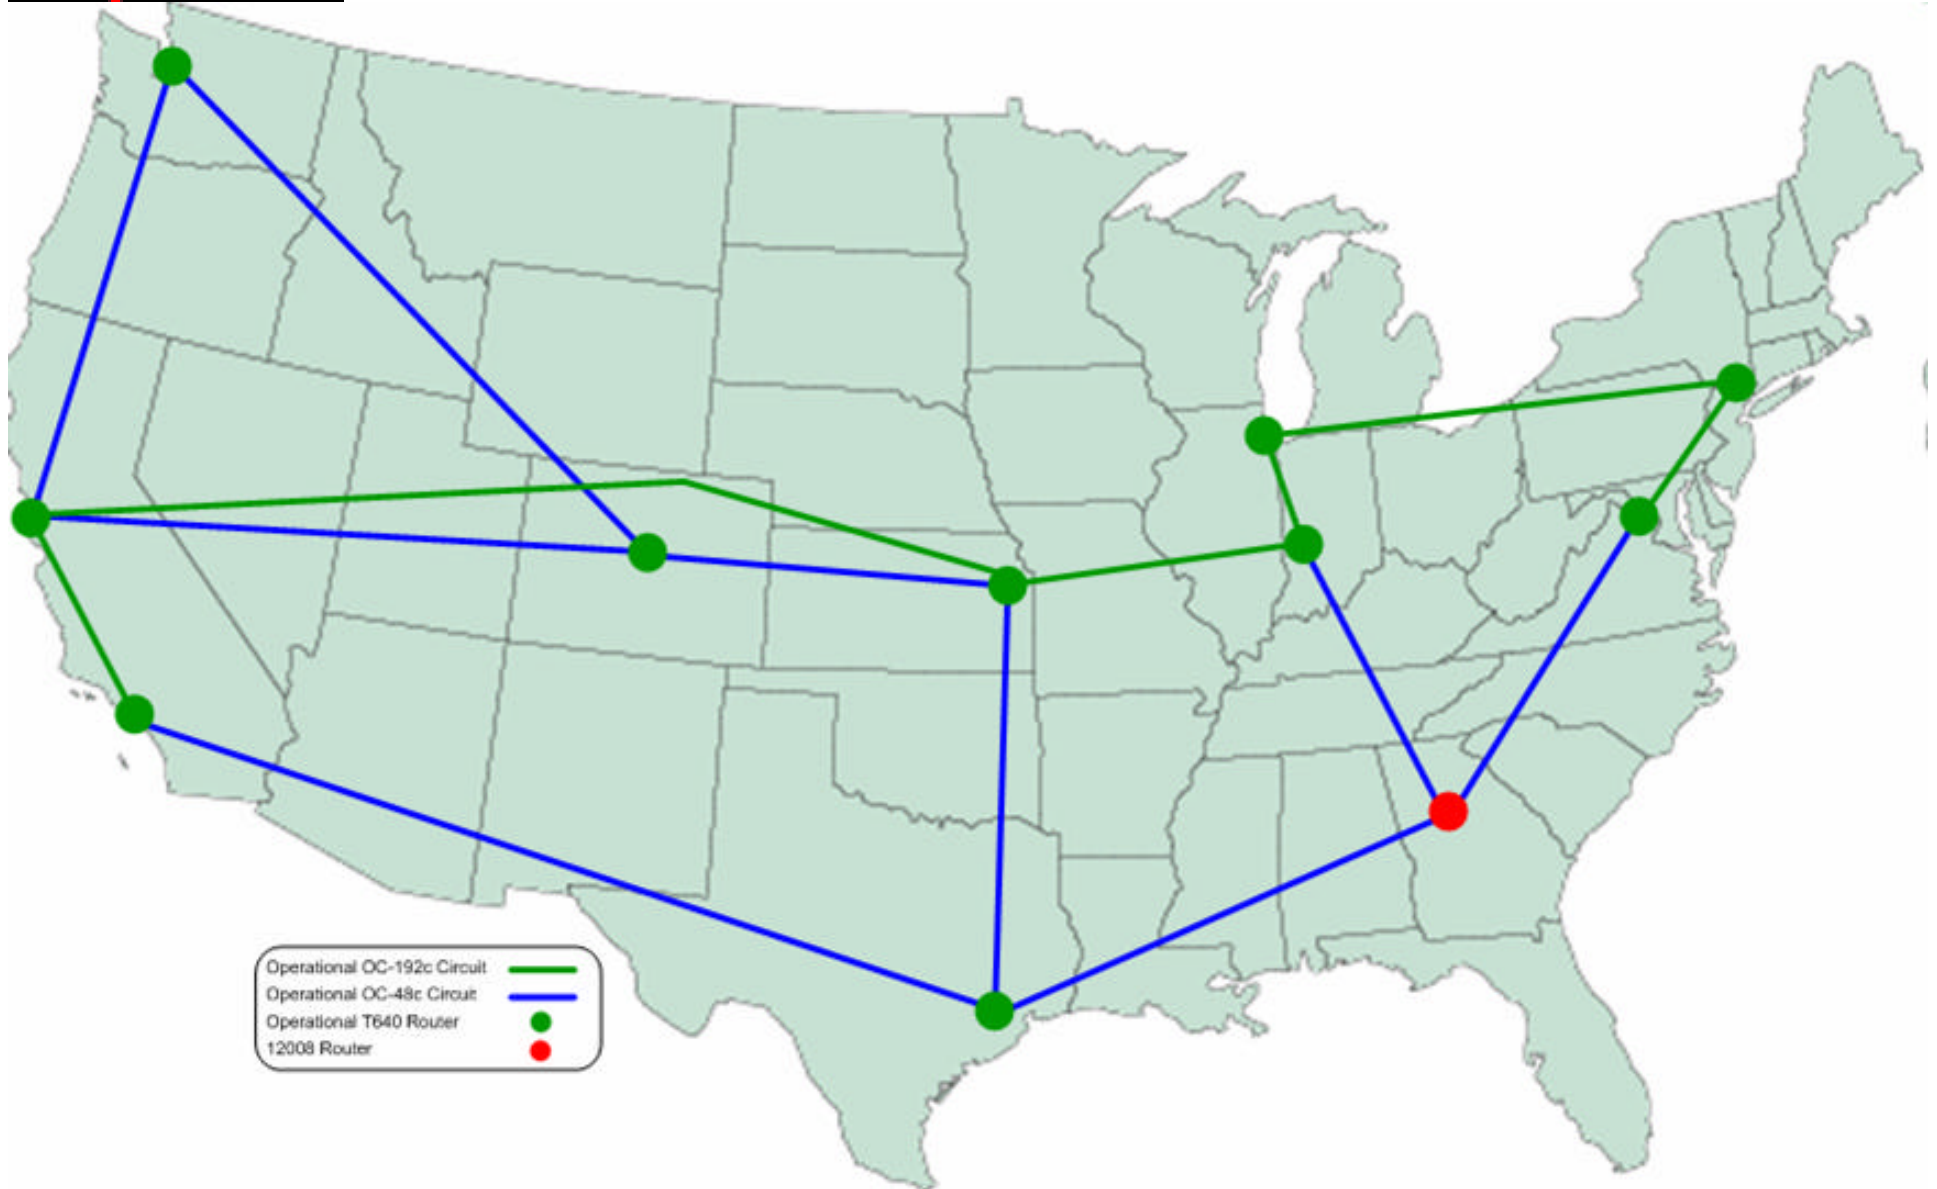
CalREN-2



The California Research and Education Network (CalREN) connects Stanford and SLAC to UC sites and provides Connectivity to the Internet2 Networks.



Abilene Topology


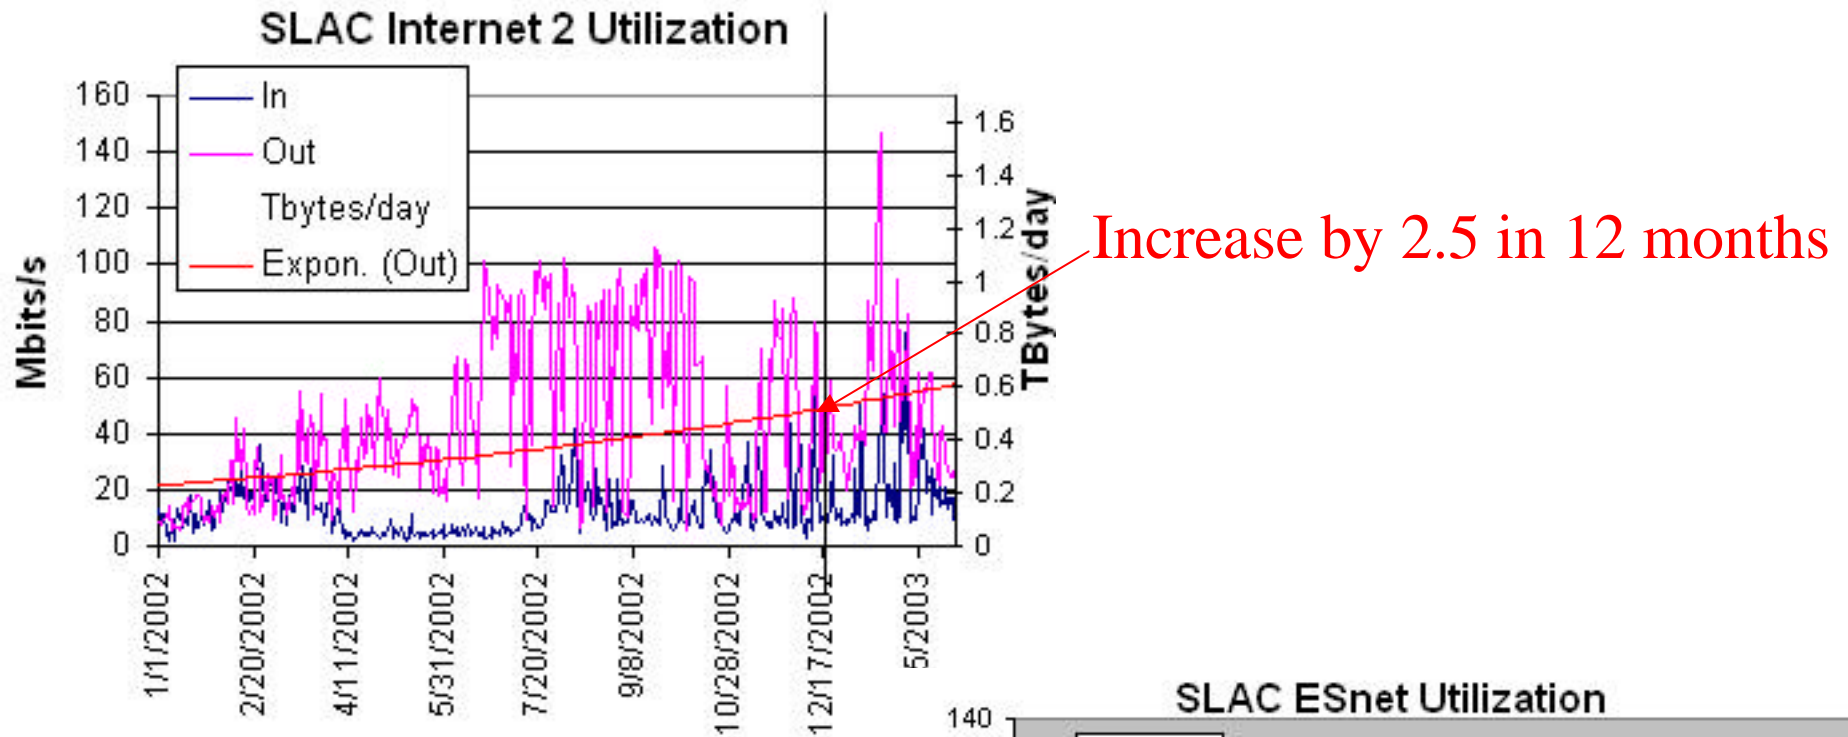


Performance Monitoring

- PingER – low impact RTT, loss etc by ping
- IEPM-BW – higher impact throughput
- Netflow – passive, characterization of traffic

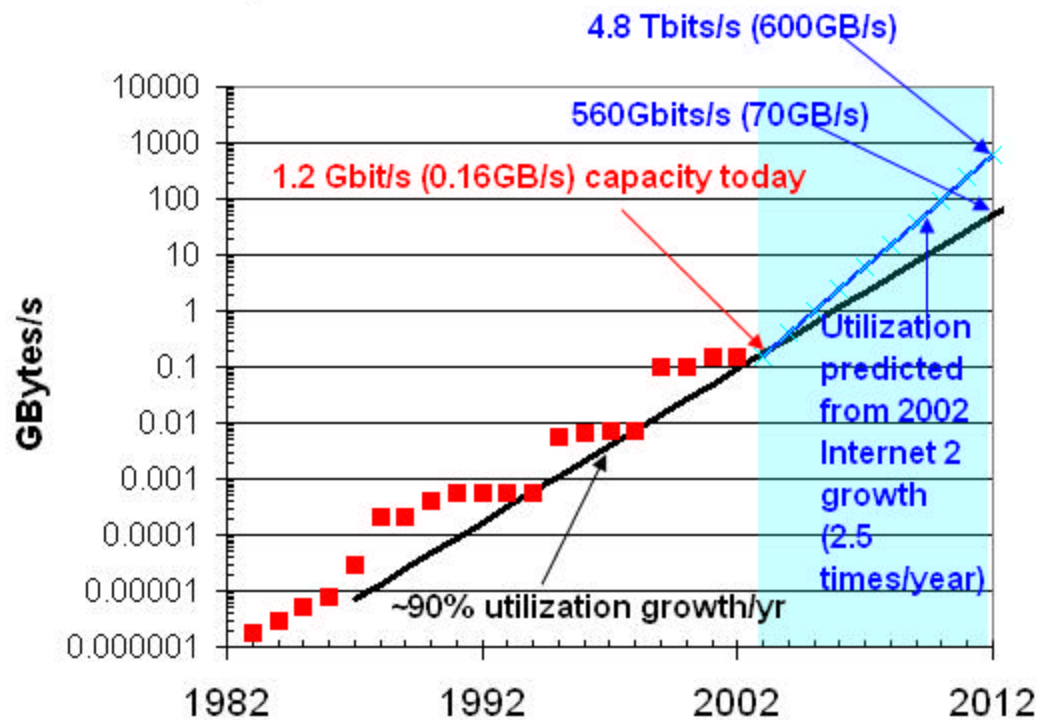


SLAC Internet Utilization 2002 -

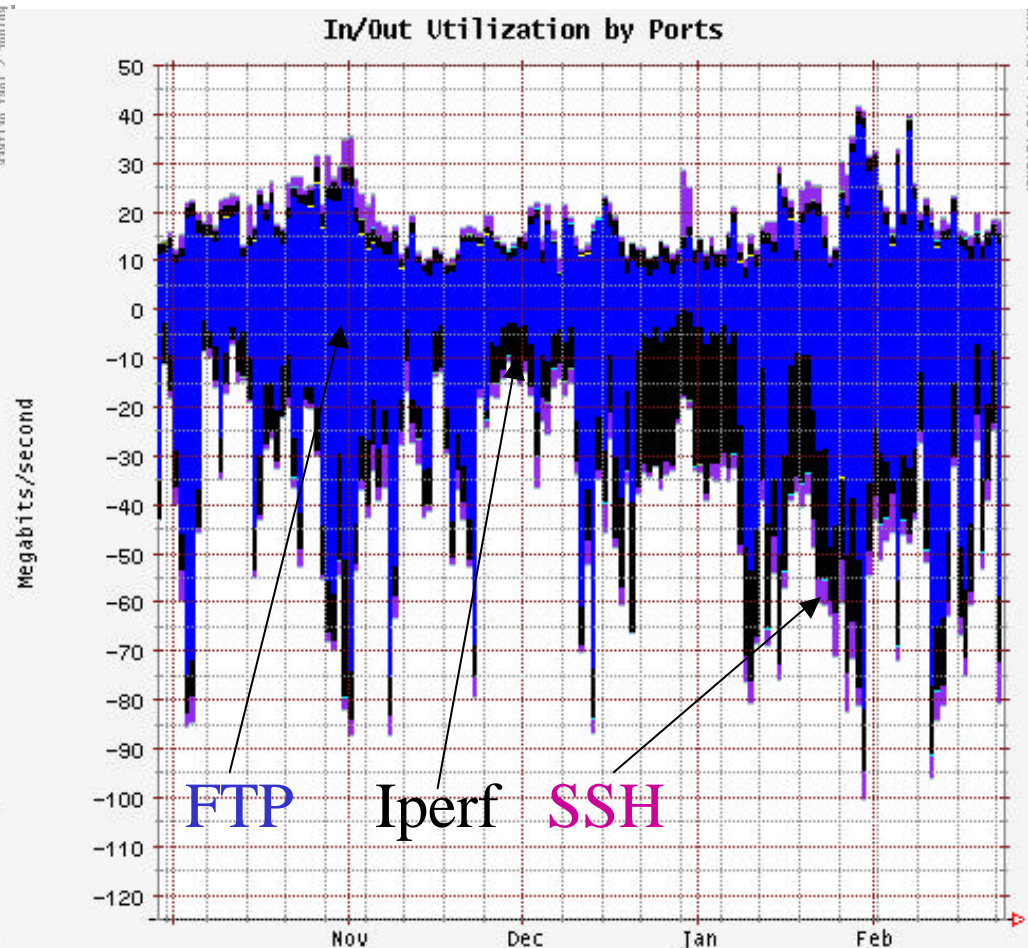
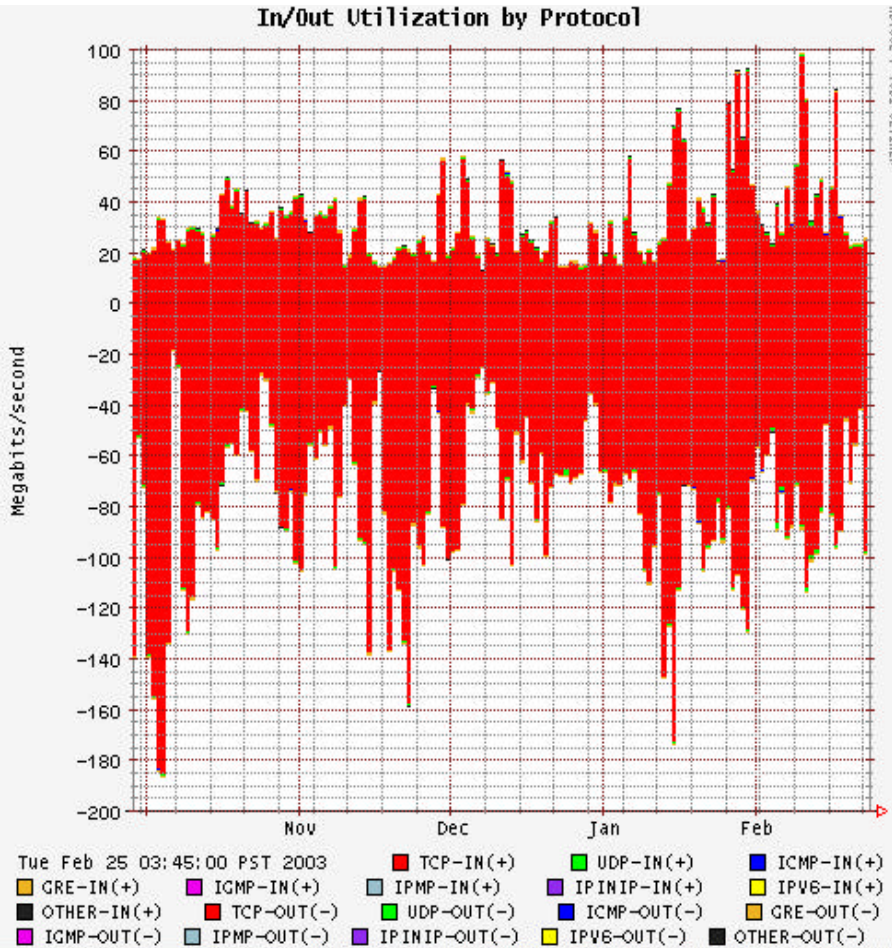


Utilization History


SLAC off-site production network bandwidth



SLAC Link Utilization (from NetFlow)



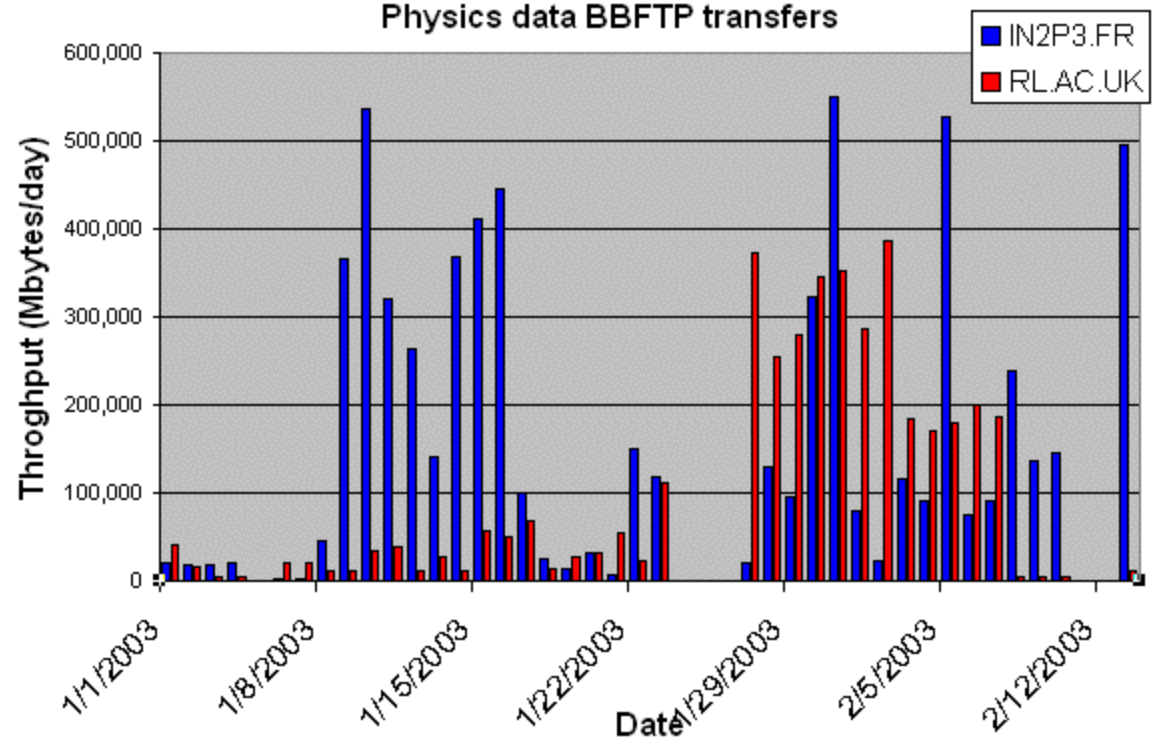
Top domains



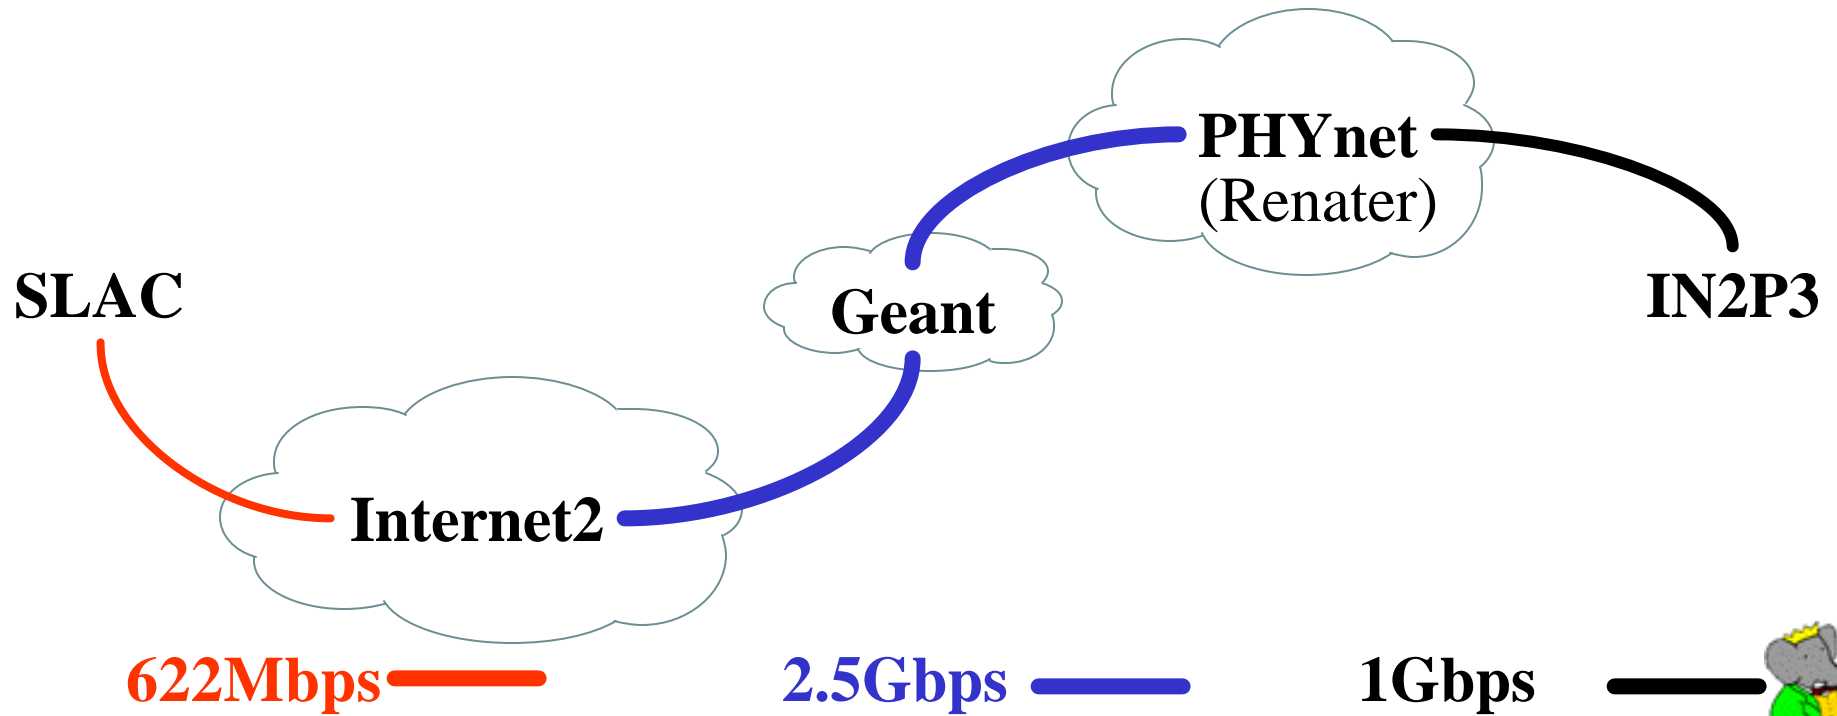

Collaborator
sites

Date

Physics data BBFTP transfers




SLAC and IN2P3



SLAC and RAL

Last Six Months (1 point per day)(data)



1Gbps RAL

TVN

JAnet

SLAC

ESnet

Geant

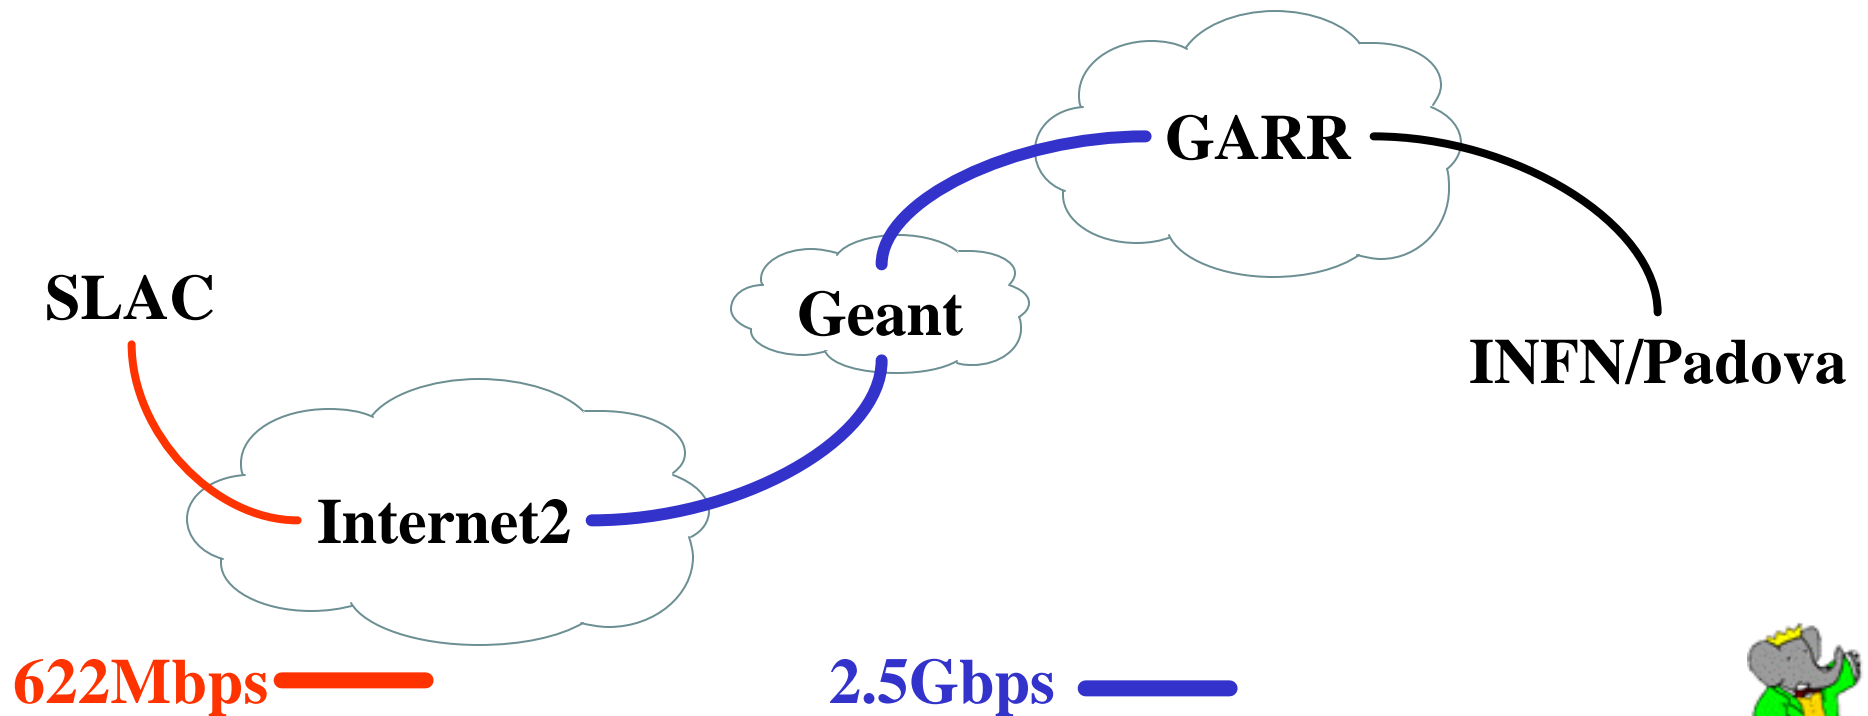

622Mbps

2.5Gbps



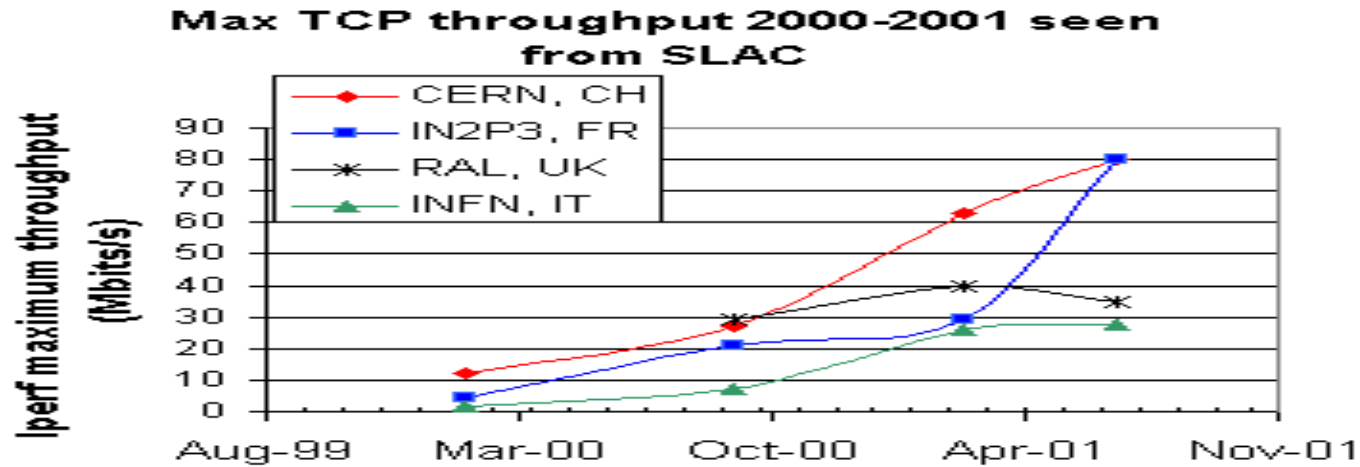
SLAC and Padova

Last 28 Days [\(data\)](#)





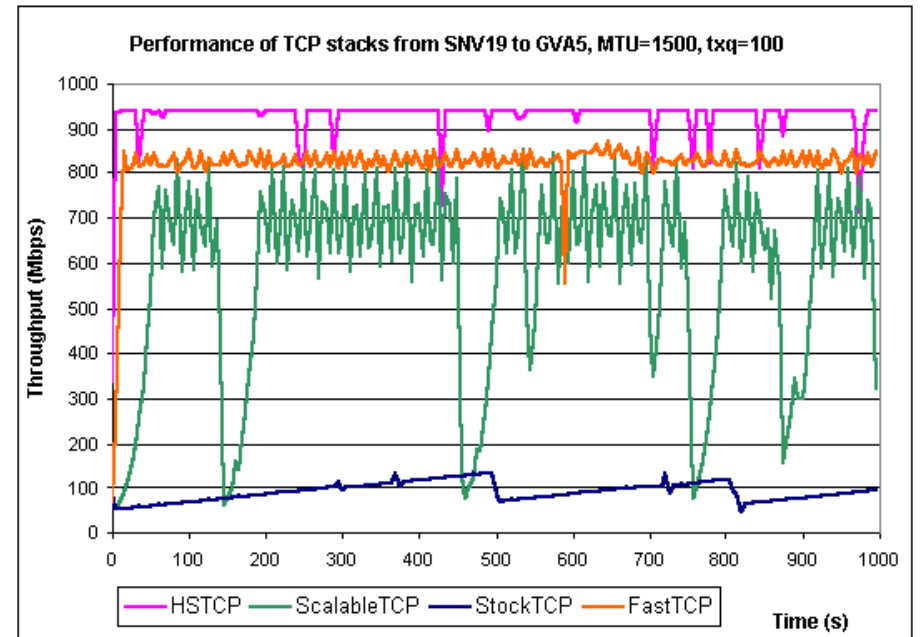


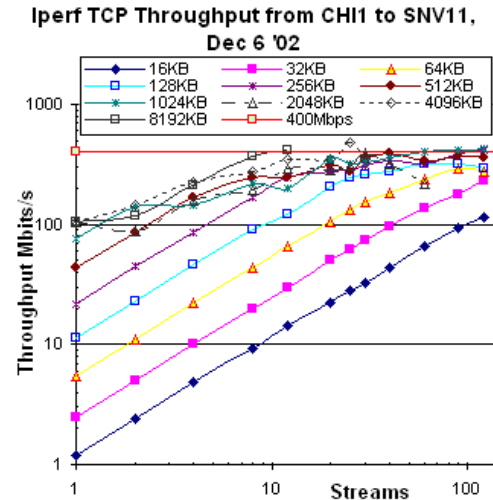
HENP is Learning How to Use Gbps Networks Fully: Factor of 100 Gain in Maximum Sustainable Thruput in 15 Months, On Some US+TransAtlantic Routes



- ◆ 9/01 105 Mbps 30 Streams: SLAC-IN2P3; 102 Mbps 1 Stream CIT-CERN
- ◆ 11/5/01 125 Mbps in One Stream (modified kernel): CIT-CERN
- ◆ 1/09/02 190 Mbps for One stream shared on 2 155 Mbps links
- ◆ 3/11/02 120 Mbps Disk-to-Disk with One Stream on 155 Mbps link (Chicago-CERN)
- ◆ 5/20/02 450-600 Mbps SLAC-Manchester on OC12 with ~100 Streams
- ◆ 6/1/02 290 Mbps Chicago-CERN One Stream on OC12 (mod. Kernel)
- ◆ 9/02 850, 1350, 1900 Mbps Chicago-CERN 1,2,3 GbE Streams, 2.5G Link
- ◆ 11-12/02 FAST TCP: 940 Mbps in 1 Stream California-CERN; 9.4 Gbps in 10 Flows California-Chicago

Improving Performance

- More bandwidth, less cross traffic
- Tune TCP/IP parameters
 - Windows, streams, MTU
- New TCP stacks
 - Tom Kelly's Scalable TCP
 - Caltech's FAST TCP
 - Sally Floyd's High-Speed TC
- Can achieve ~1Gbps trans-Atlantic
 - SLAC and collaborators hold the world record (Nov-02)



923Mbps for 6800miles



Other gotchas

- Linux memory leak
- Linux TCP configuration caching
- What is the window size actually used/reported
- 32 bit counters in iperf and routers wrap, need latest releases with 64bit counters
- Effects of number of packets in queue for NIC (txqueuelen)
- Routers do not pass jumbos
- Some applications have problems with big windows
- Firewalls



Toward a Monitoring Infrastructure

- A ubiquitous network monitoring infrastructure
 - Grid Middleware
 - Guide Applications
 - Troubleshooting
 - Resource Allocation
 - Set expectations



Summary

- Multiple high-speed connections
 - Optimize routing
 - Backup paths
- Well Engineered Networks
 - Negligible packet loss, Good Round Trip Times
 - Traffic doubled on ESnet every year since 1990
 - Need continued upgrades to keep up with demand
- Transmitting TBytes/day today
- Room for improvement (bottlenecks at edges)
 - Needs funding &
 - Improved TCP or other protocols
- TCP works up to Gbits/s rates, but needs wizards today to make it perform



More Information

- High-performance Monitoring
 - www-iepm.slac.stanford.edu/bw/
- TCP stack evaluation and 10GE testbed
 - www-iepm.slac.stanford.edu/monitoring/bulk/fast/
- Measurement reports
 - www.slac.stanford.edu/comp/net/bandwidth-tests/antonia/html/slac_wan_bw_tests.html
- Land Speed Record
 - www-iepm.slac.stanford.edu/lsr/

Any Questions ?



Throughput from SLAC to RAL between May 2002 and February 2003

