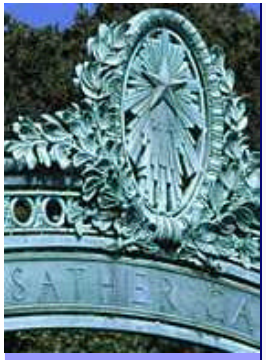

Using PacketShapers to control UC Berkeley I$P costs

ken lindahl
Communication & Network Services
UC Berkeley

lindahl@uclink.berkeley.edu

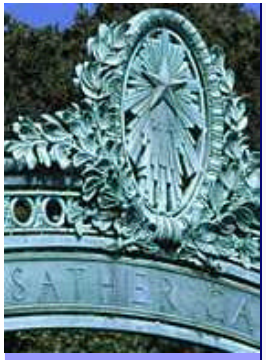

challenge: to limit only commodity transit traffic

- UCB gets commodity transit from CalREN-2; costs are usage-dependent…
- … also gets some things that we don't want to rate limit, since our costs for these are fixed, not usage-dependent:
  - Abilene
  - CUDI, NevadaNet
  - commodity peering at PAIX, LAAP, SDNAP
- we tried a few things…

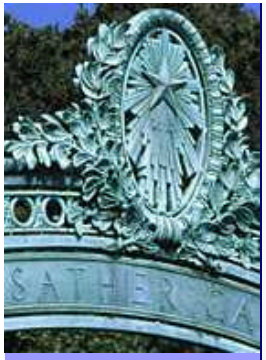

2nd try: ResHalls

- packetshaper on fast ethernet link between ResHall backbone and campus router

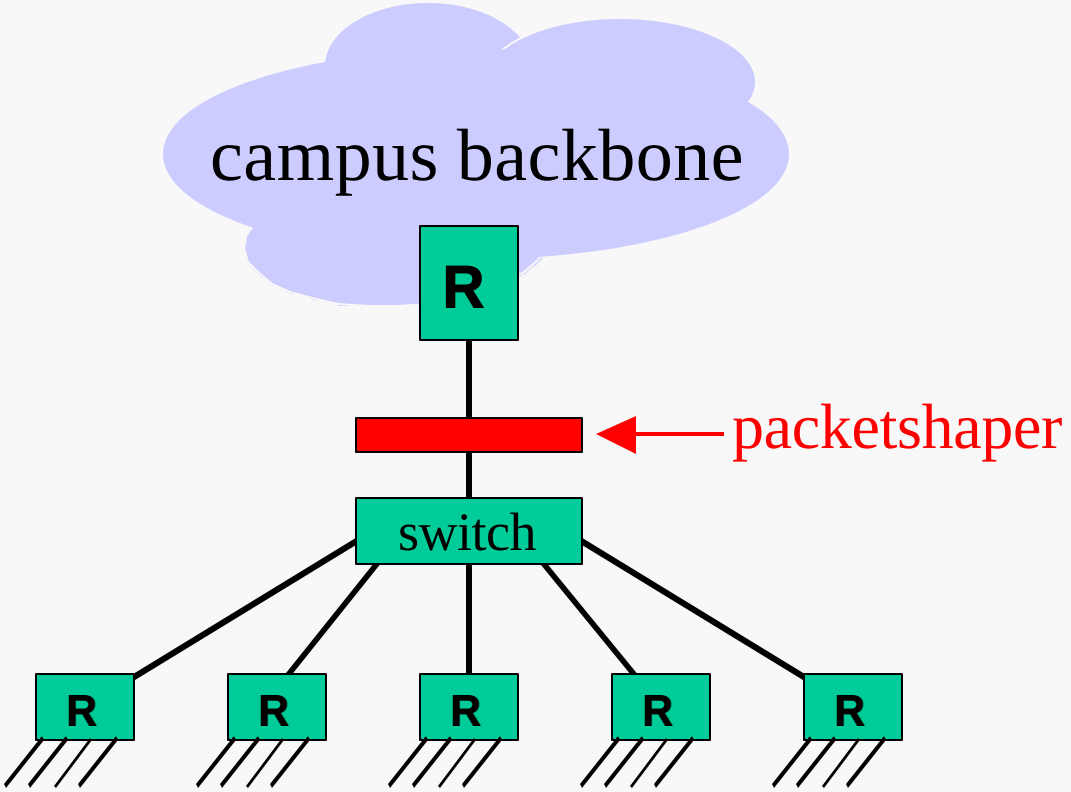

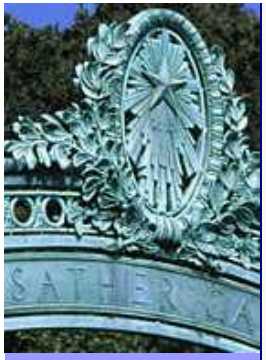

2nd try: ResHalls (2)

- analysed one week's worth of netflow records to identify the 2400 most heavily used Internet2 prefixes;
  - these accounted for 99.2% of Internet2 traffic during the week (Jan 2000).
- created {in,out}bound classes based on those prefixes;
- applied fixed partition to "Default" class.
  - so approx 99.2% of Internet2 traffic was excluded from the rate limit

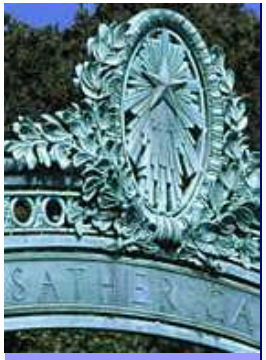

2nd try: ResHalls (3)

- it worked, but…
  - never found time to keep the exempt classes up to date;
  - the number of Abilene-announced prefixes grew over time;
  - so did the numbers of other rate-insensitive prefixes.
- so, prefix-based class definitions didn't scale well;
- and, we still had to worry about data to/from the rest of campus (ROC).

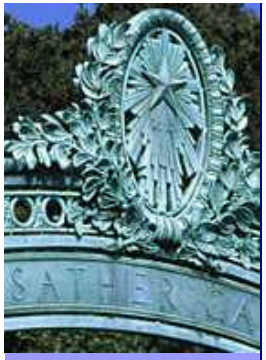

3nd try: ResHalls and ROC

- re-engineered the campus border topology so that commodity transit data followed a different path than non-transit data;
- inserted the packeteer into the commodity transit-only path.
  - requires a separate connection to CalREN-2 for commodity transit traffic.
  - requires two packetshapers since total b/w is greater than 100 Mbps.

3nd try: ResHalls and ROC (2)

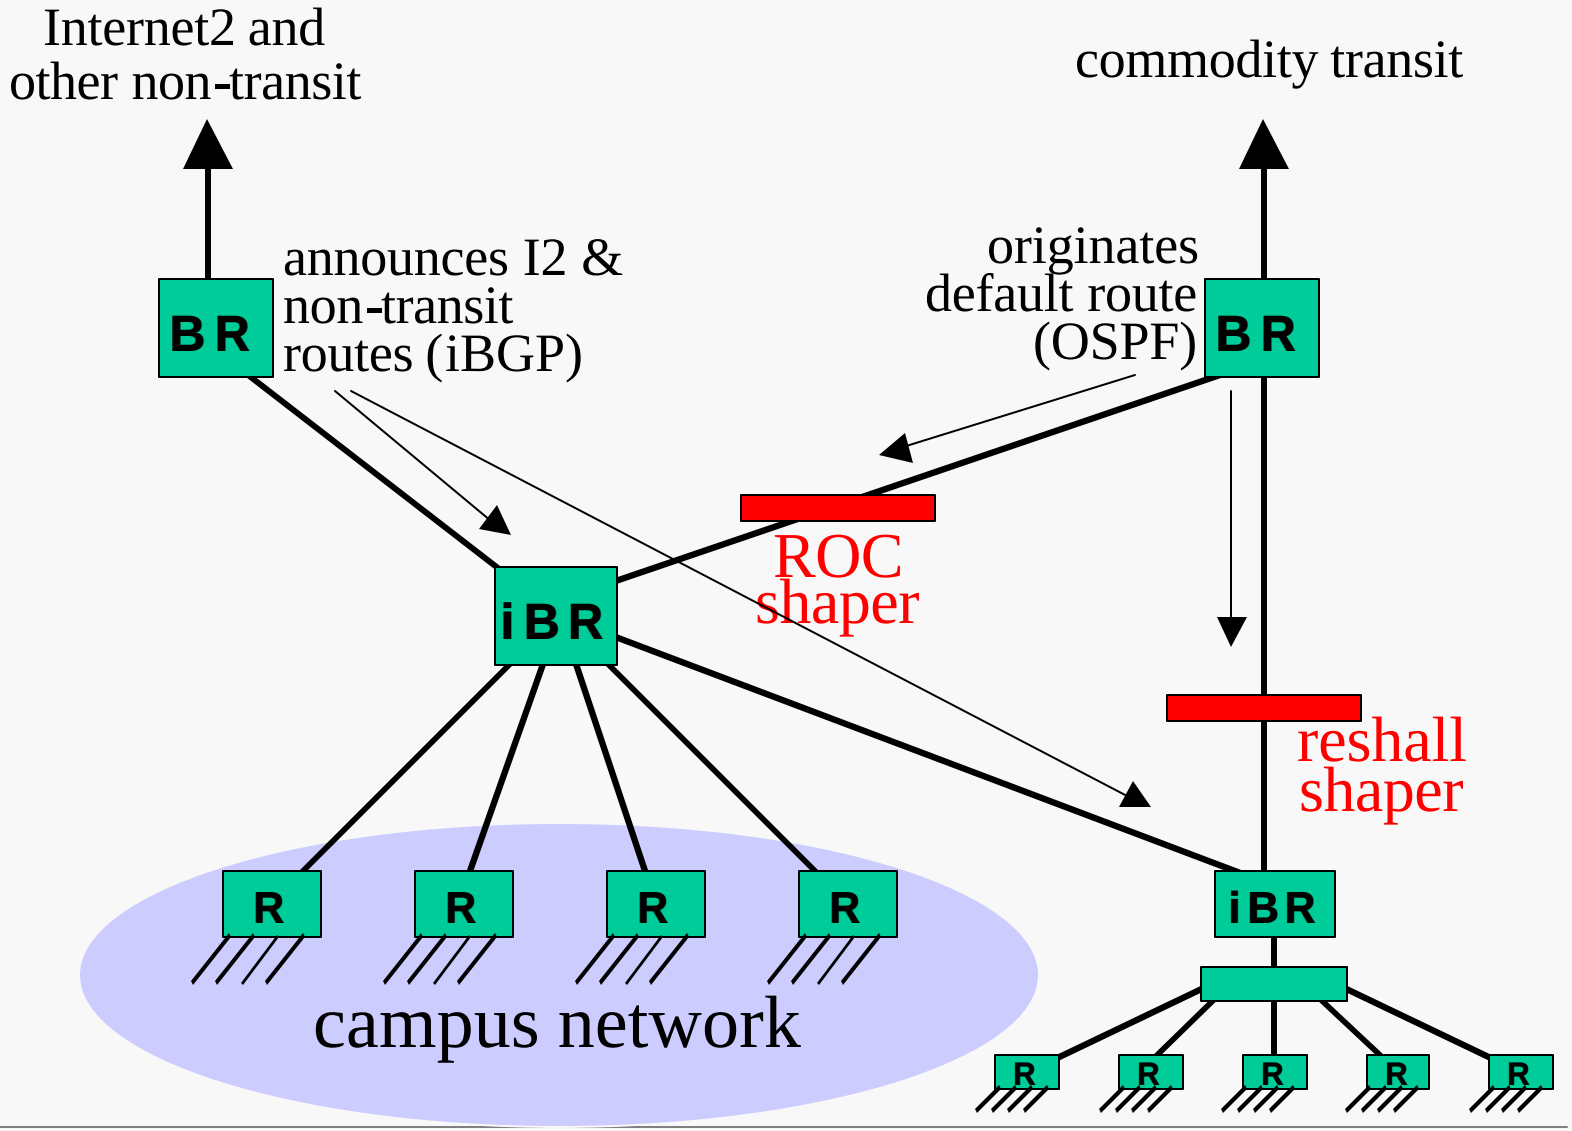

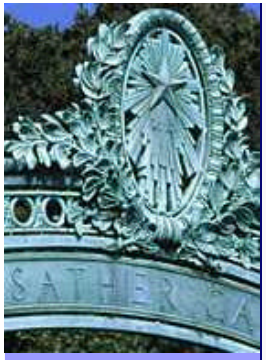

3nd try: ResHalls and ROC (3)

- extremely simple configuration:
  fixed-size partitions on /Inbound and /Outbound classes
  - no complex class definitions requiring maintenance;
  - conserves packetshaper resources for other matters, such as passive analysis of “interesting” patterns;
  - very simple to explain to campus users;
  - not concerned about applications hiding behind non-standard ports;
  - avoids all concerns about deciding which packets are “better” than others.

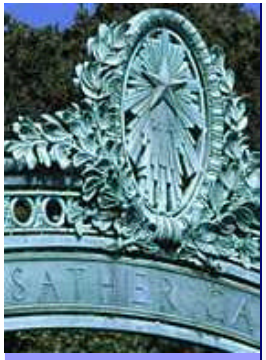

result: ResHall traffic

- rate-limited at 40 Mbps (for which they pay)

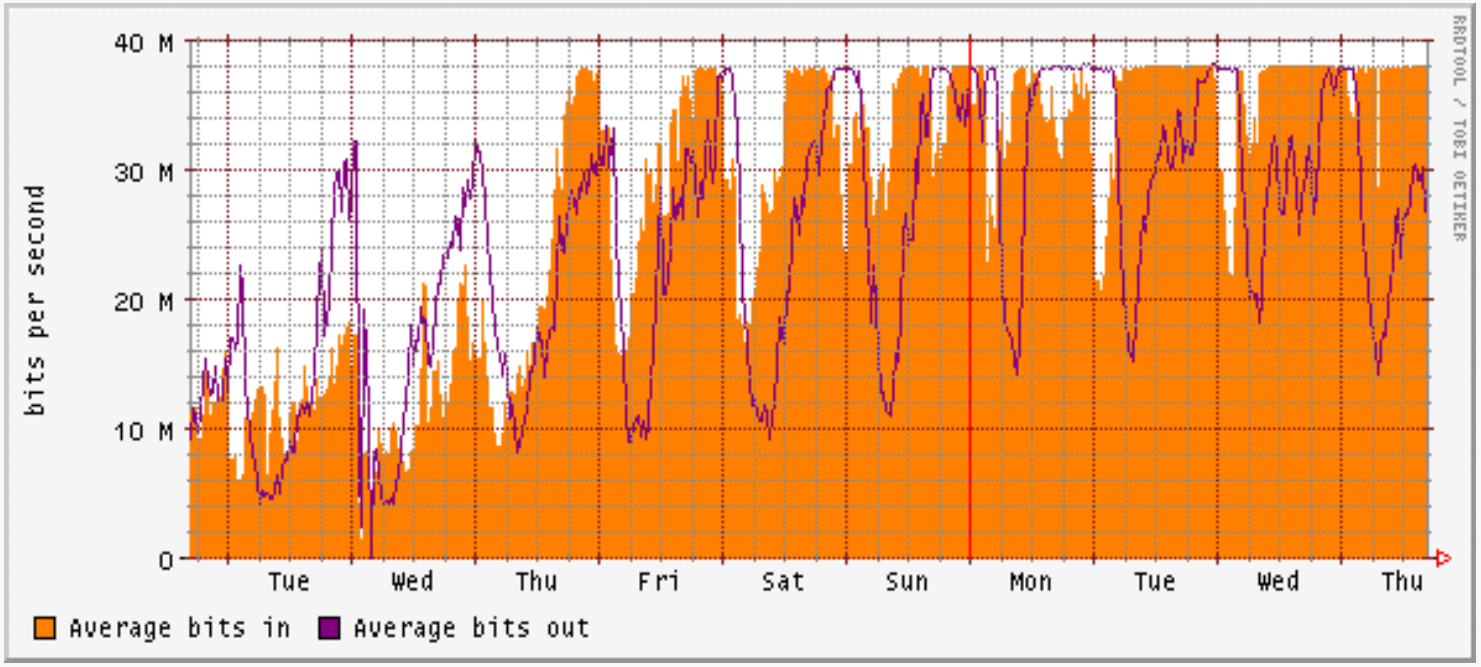

  - fast ethernet link to ResHall backbone (15-24 Jan 2002)

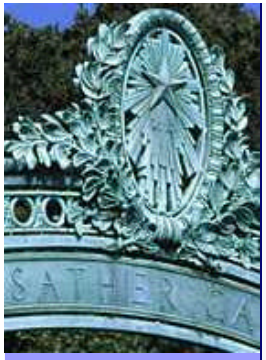

result: ROC traffic

- Rate-limited at 70 Mbps

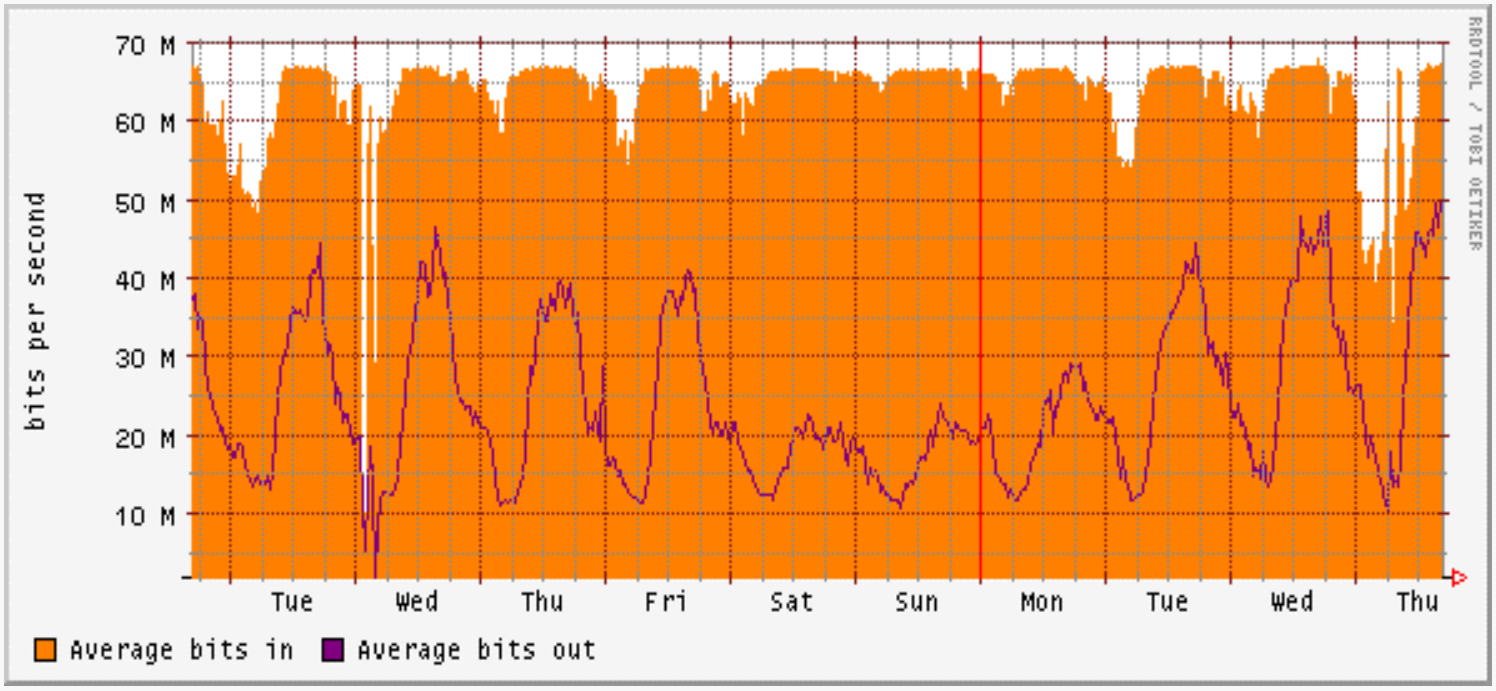

  - fast ethernet link to ROC (15-24 Jan 2002)

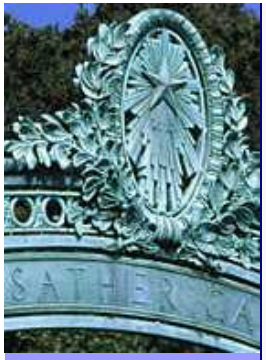

result: I$P traffic

- Total commodity transit traffic:

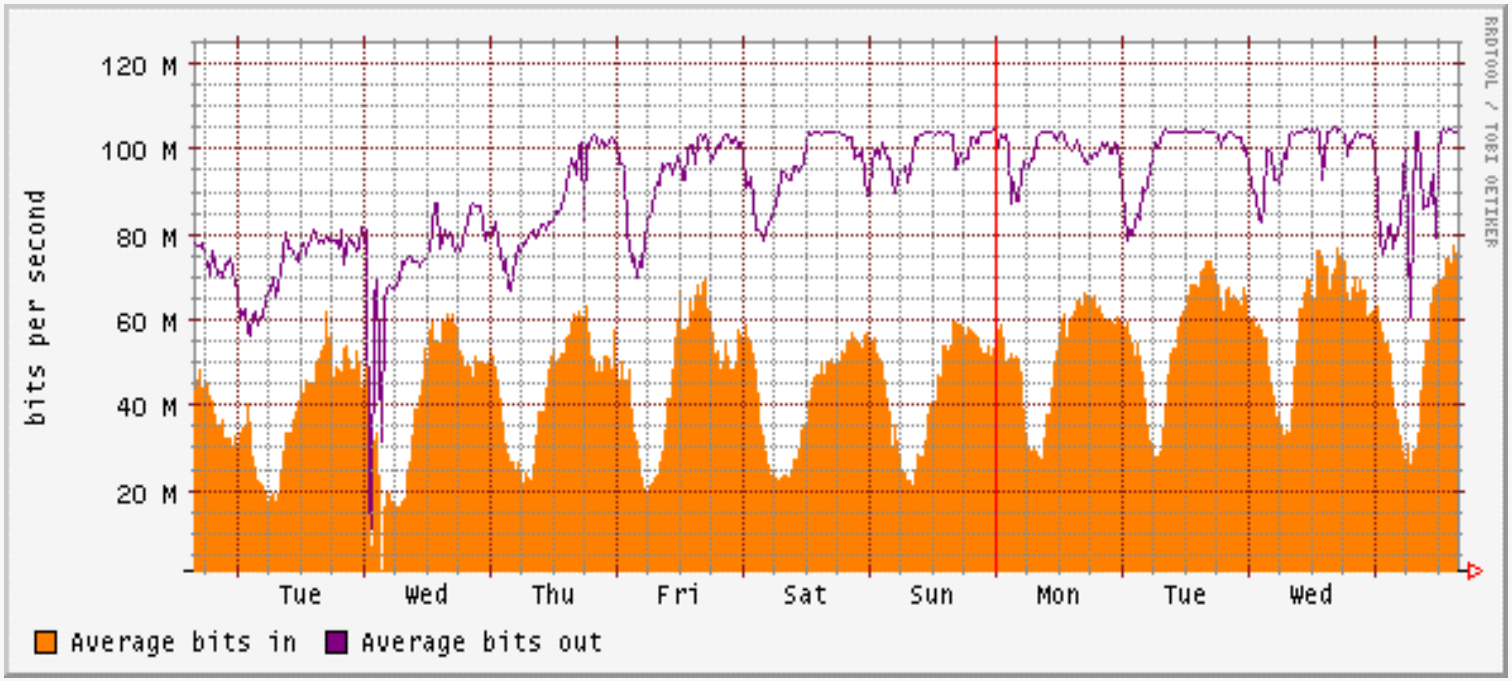

  - OC-3 ATM link to CalREN-2 commodity transit service (15-24 Jan 2002)

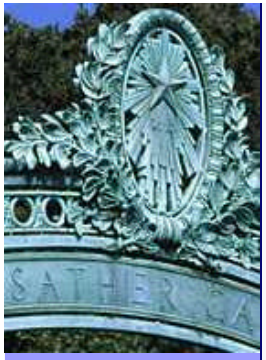

the future

- (legitimate) campus demand for commodity transit is increasing; expect greater than 100 Mbps for ROC in Fall 2002.
  - will need next generation packetshaper capable of handling the higher data rate, and with GbE ports.
  - campus must develop guidelines for growth in commodity transit traffic, and allocate funds to support that growth.
  - need guidelines for when a group is consuming too much of the common b/w, and should pay for additional capacity.

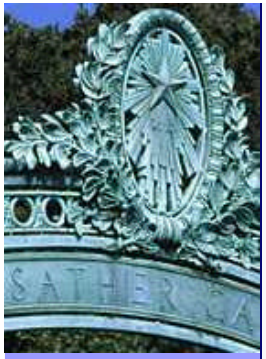

summary

"You can't go fast only on Internet2" [1]

but you can control how fast your commodity traffic goes.

[1] Joe St Sauver

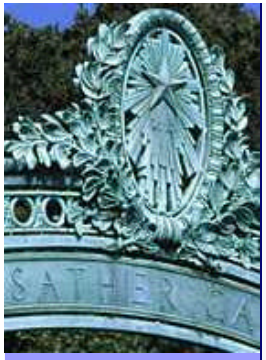

epilog: 23 Jan 2002

- latency across the ROC packetshaper link was around 400+ ms (normally 1-2 ms)
- discovered that SETI@home traffic was competing too well for b/w in the ROC 70 Mbps partition.
  - had previously applied a 0-30 Mbps burstable partition on SETI@home, to “cap” spikes from the server.
  - SETI folks were not experiencing the horrid latencies.
- removed the SETI@home 0-30 Mbps partition at 7:00PM 23 Jan 2002.

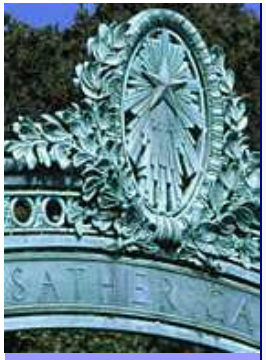

epilog: 24 Jan 2002

- that same night, the server went offline at midnight; came back online at 7:30AM 24 Jan 2002.
- server spiked to 50-60 Mbps, causing 1200-1500ms latencies on the packetshaper link.
- SETI folks suggested giving their traffic a lower priority. (class gets priority 2)
- seems to be working well: they consume otherwise unused b/w; rarely drop below 10 Mbps and get 35+ Mbps in the early morning:

SETI@home
before 24 Jan 2002

- 22-28 Mbps, inside the ROC 70 Mbps limit

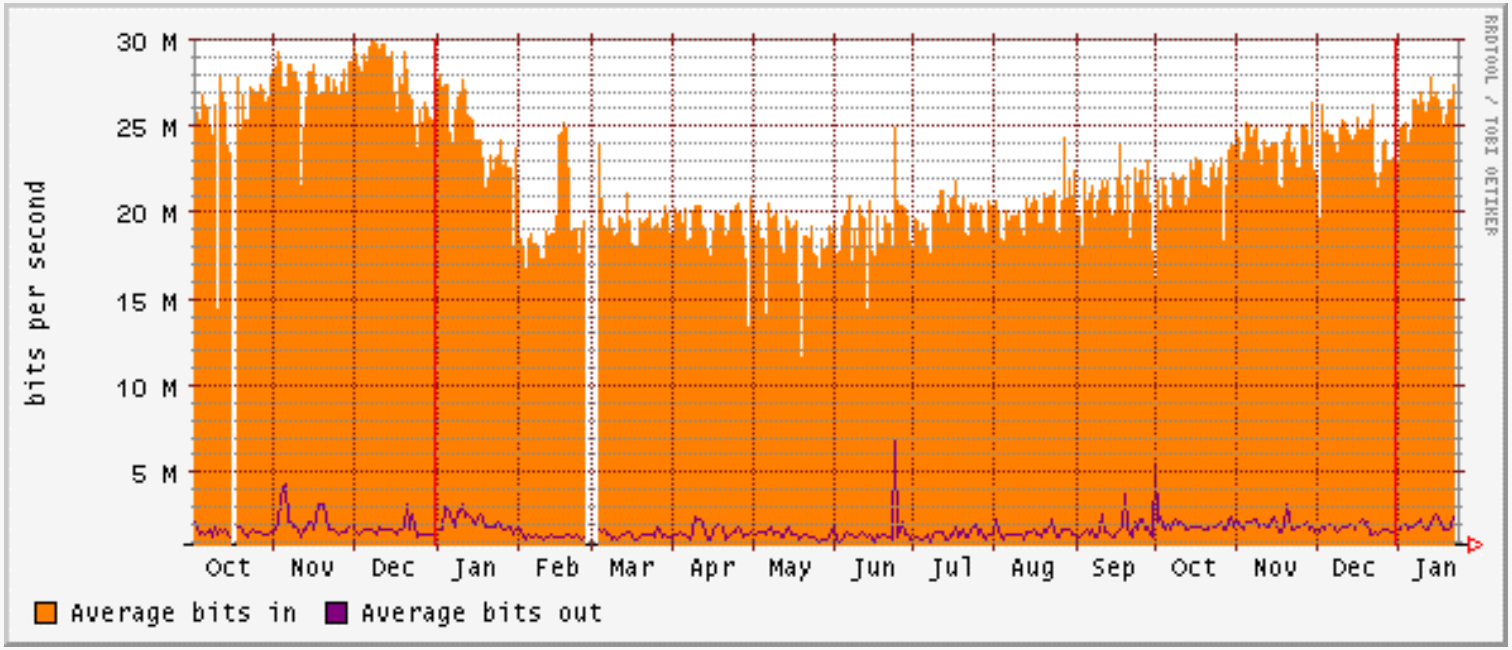

  - fast ethernet link to Space Sciences Lab (Oct 2000 – Jan 2002)

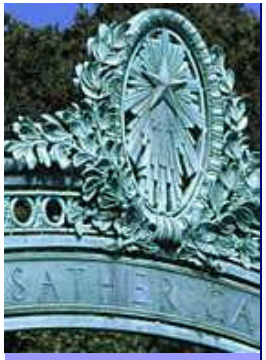

SETI@home
after 23 Jan 2002

- inside the ROC 70 Mbps limit

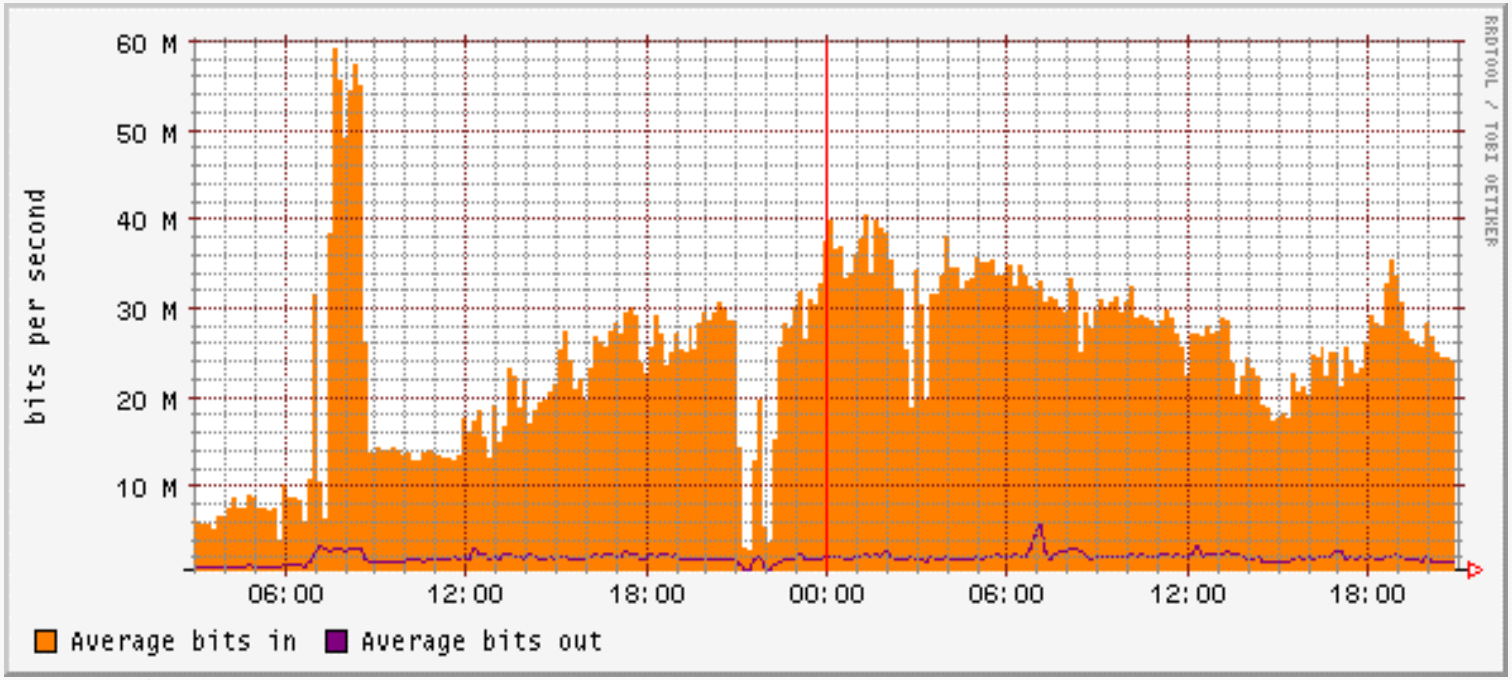

  - fast ethernet link to Space Sciences Lab (Oct 2000 – Jan 2002)