Campus Bandwidth Management: Approaches and Tradeoffs

This is a work in progress. Oct 31, 2003

Approach	Advantages	Disadvantages	Examples
Do Nothing	• Simple	• Unfair • Expensive • Mis-match between usage and cost recovery, especially severe if university is charged per-bit, but performs cost recovery by charging flat fees • Mission of university may be impeded by inappropriate use	*Many*
Per-IP Quotas (Rate-Based)	• Arguably "fair" • Can tune quotas so that conforming traffic rarely experiences congestion • No need for application-level classification • End-system portability is supported (since all ResHall IP addresses are policed identically)	• IP addresses become an artificially rare commodity (consider impact on IPv6) • Additional router complexity • May impede deployment of meritorious high-bandwidth applications (especially if limits apply to Internet2 traffic) • Inability to burst once in a while	*U. Penn* An overall rate limit is applied to outbound ResHall traffic. Additionally, rate-limiters (one per IP address) are installed on the edge router and applied only to outbound traffic. [talk] [updated talk]

Approach	Advantages	Disadvantages	Examples
Per-IP Quotas (Volume-Based)	• Top talkers can be isolated by placing them in a penalty box • Negative feedback loop encourages users to modify their own behavior • No need for application-level classification • Ability to burst once in a while	• IP addresses become an artificially rare commodity (consider impact on IPv6) • May impede deployment of meritorious high-bandwidth applications (especially if limits apply to Internet2 traffic) • Additional router complexity • Additional accounting complexity • Usage and penalty status need to be communicated quickly to average users	*North Dakota State University* Quotas apply only to ResHall users. Quota is 300 MB per day per user. Users who exceed their quota are placed in a shared pool rate-limited to 256kbps. [talk] [ResNet] *University of Waterloo* Residence hall users subjected to per-user quotas of the form "*x* MB in last *y* days". In addition the residence hall traffic aggregate is given a guaranteed minimum share of external bandwidth through CB-WFQ. [more info] *Iowa State* Residence hall users who exceed a specific level (currently 200 MB), are transferred to a "slower Internet connection". As abuse continues, offending users are shifted to ever more restricted traffic classes. User quotas are reset at the end of each day,

Approach	Advantages	Disadvantages	Examples
			except for those in the rate-limited classes, for whom a 24-hour moving average is applied to determine when they are returned to a less restrictive traffic class. [more info] *Virginia Tech* see below
Per-Class Quotas (Rate-Based)	• Can balance use among different user communities • Can tune so that conforming or exempt classes rarely experience congestion • Easy to implement (if not discriminating between commodity and Internet2 traffic) • No need for application-level classification	• No fairness within classes • May impede deployment of meritorious high-bandwidth applications (especially if limits apply to Internet2 traffic)	*UC Berkeley* Packeteers in front of a campus edge router separately rate-limit commodity traffic to/from residence halls and to/from the rest of campus (ROC) traffic. Two PacketShapers are required because the total bandwidth exceeds the 100 Mbps. Routing has been engineered to keep ResHall and ROC traffic separate. [talk][1] *Virginia Tech* Complex hybrid approach that primarily employs class-based policing, but also makes use of application-based policing and a penalty box scheme. Off-

Approach	Advantages	Disadvantages	Examples
			campus traffic from residence hall subnets is policed to 60 Mbps aggregate and off-campus traffic from the campus news server is policed to 5 Mbps. "Nuisance applications" are policed to 10 Mbps in aggregate (profiles are generated manually). Finally, individual users are placed in one of three classes: Class 0 (unpoliced), Class 1 (policed to 1.5 Mbps), and Class 3 (policed to 250 Kbps). When users exceed a certain threshold (currently 650 MB) in a 24hr period, their class is incremented; if they stay under threshold, their class is decremented. (The CB-WFQ scheme described in the talk below is not currently in use.) [talk] *University of Washington* Total network bandwidth from the residence halls to off-campus commodity destinations is limited to 100 Mbps. Off-campus access to

Approach	Advantages	Disadvantages	Examples
			common server ports (Web, FTP, IRC, etc) in the residence halls is blocked. Inbound peer-to-peer traffic is rate-limited to 20 Mbps; outbound peer-to-peer traffic is limited to 2 Mbps. [residence hall computing policy] *UC Santa Cruz* see below
Per-Class Proportional Sharing	• Restricted traffic classes can use unused capacity	• No fairness within classes • May impede deployment of meritorious high-bandwidth applications (especially if limits apply to Internet2 traffic)	*University of Waterloo* Residence hall traffic is given a guaranteed minimum share of external bandwidth through CB-WFQ. (see above) *Texas A&M* Planning to support four application classes. Per-session admission to classes. Diff-serv edge marking, policing, and stateless core queueing. (Currently using per-application rate-limits.) [talk]

Approach	Advantages	Disadvantages	Examples
Per-IP Proportional Sharing	• Arguably "fair" • No surprises (users get the service they pay for) • [additional praise]	• IP addresses become an artificially rare commodity (consider impact on IPv6) • May impede deployment of meritorious high-bandwidth applications (especially if limits apply to Internet2 traffic) • Additional router complexity • Many queues required • Care must be taken not to restrict Internet2 performance	*No known deployment examples*
Usage-based Charges After Threshold	• Economically rational (users who get the most value from a scarce resource pay the most for it) • Fair • Negative feedback loop for heavy users • Can be tuned so that most users pay flat monthly rate; similar to pricing of department printers for students, of cell phones, etc. • [additional praise]	• Additional accounting and billing complexity • Need system to collect usage stats (*e.g.* NetFlow)	*Cornell* Planning to charge each department a monthly fee that includes a WAN usage component. Rate structure to include a mix of port fees, infrastructure tax, and usage fees. Per-megabit usage fees will only kick in for use above a certain threshold (adjusted so that 80% of IP addresses will avoid usage fees). Monthly bills to the departments will include enough detail to support recursive usage-based charges to individual users or

Approach	Advantages	Disadvantages	Examples
			research groups. NetFlow-based billing system using Apogee software and home-brewed scripts. [white paper] [web site] *University of Kansas* Applying artificially low usage based charge to ResHall users. Only heavy users will feel the usage based fees; ordinary users will be charged a flat rate.
Per-Application Quotas (Rate-Based)	• Majority of problems often caused by small number of applications • Tool to reduce illegal use of network (*e.g.* illegal distribution of copyrighted materials) • "Magic bullet" middlebox • Automatic maintenance through "bad apps du jour" subscriptions • [additional praise]	• Must pass judgment on which applications are "good" and which are "bad" • Performance impact (QoS appliances are designed to handle a scare resource and therefore generally lag routers in their ability to handle high speeds or maintain very low loss rates for "good" traffic) • Loss of transparency (e.g. rewriting of TCP window size) • Complex and dynamic configurations complicate performance debugging	*UC Santa Cruz* Allot NetEnforcer deployed between ResNet and commodity/Internet2 access link. Traffic is classified into four priority levels: High (web, ssh), Medium (everything except peer-to-peer), Low (peer-to-peer), Blocked (worms). [talk][2] *Virginia Tech* see above *University of Washington* see above

Approach	Advantages	Disadvantages	Examples
		• Application profiling creates a cat and mouse game that the mouse will win (*e.g.* http, https, proxies, random port numbers, ssh, etc.) • [additional criticism]	
Outsource Residential Networking			*University of New Mexico*
Block Servers (with NAT or firewall)	• Can apply only in "bad neighborhoods" (*e.g.* residence halls)	• Destroying end-to-end transparency can restrict deployment of numerous advanced applications (*e.g.* VoIP, research-oriented peer-to-peer) • Potentially sever performance impacts • Motivated users will learn to punch through	*We know you are out there!*

Footnotes

1. Talk addenda (10/25/2002): ResHall rate limit is 60 Mbps in each direction and ROC rate limit is 100 Mbps in each direction; SETI@Home has purchased its own ISP service and is no longer in Berkeley's IP address space
2. Talk addendum (10/25/2002): UCSC has acquired a faster Allot box with more memory; they are still experiencing some problems with interactive performance.