7
Case Studies

You may have noticed that there is not a single definitive set of instructions on "how to fix your slow Internet connection" anywhere in this book. This is because every environment has different needs and resources, and the techniques used for bandwidth management will vary greatly depending on your local environment. To put this into perspective, here are some real life examples of how the methods and techniques presented in this book have solved real world bandwidth management problems.

KENET, Kenya

Established in 1999, The Kenya National Education Network (KENET) connects educational institutions and research centres with the goal of broadcasting knowledge throughout the country. Currently, there are 13 members directly connected to the main node in Nairobi and 40 additional members participating in the network via Kenya's backbone Telco. The main node consists of a 3 Mbps uplink via a leased line, and a 3 Mbps downlink via VSAT.

Clearly, available bandwidth is limited. Members typically range from 64Kbps to 960 Kbps of bandwidth usage, many of whom connect to KENET via leased lines from the local Telco. These leased lines then terminate at KENET on E1 lines with 2 Mbps capacity. In addition, some of the members have their own VSAT downlinks and only uplink via KENET.

Most sites do not have skilled staff managing the network, and are often unaware of the dangers and costs involved when bandwidth is mismanaged.

Problems

Initially, certain networks were set up improperly due to a lack of trained staff. For example, one network consisting of 20 machines was operating without a caching proxy server. Servers and routers were often poorly configured. The major consequence of this neglect was that the bandwidth became widely mismanaged – usually clogged with peer-to-peer traffic, spam, and viruses.

In turn, this led to KENET IP addresses being blacklisted on the Internet, making services and sites completely unreachable. The already low capacity bandwidth became completely unusable for most institutions.

Analysis

An analysis of each institution's traffic behaviour was made using MRTG (traffic grapher) and FlowC (graphical packet analyser). This established a baseline so that changes to the network could be monitored for their effects. In addition to investigating the problematic institutions' network configurations, the software in use at each site was examined.

The minimum skill level required of staff at participating institutions was determined, and the possibility of replacing some software with FOSS (Free and Open Source Software) was researched.

Solutions

A regular training program for staff was started in 2004. Training topics included network management, security, and monitoring using FOSS tools. Training sessions also educated staff to interpret traffic graphs and detect anomalies.

On the networks, aggressive Access Control Lists (ACLs) were installed on routers to restrict access to include only approved services. Custom servers were set-up by KENET to address specific problems at each institution. While the configuration details varied between locations, a decision was made to standardise the network on a uniform platform: Firewalled FreeBSD. This solution proved to be stable, secure, reliable, and FREE! Uniform solutions were chosen for each software component as well: MTA (Qmail), Spam Filter (Spamassassin), Mail Antivirus (ClamAV), Proxy (Squid), and BSD traffic shaping.

Here are the results of applying good bandwidth management techniques to three sites within KENET.

Site One: firewall & proxy server

The institution at Site One had 960 Kbps of available bandwidth. This institution's bandwidth needed to accommodate approximately 300 computers connected to Internet, plus two Mail and two DNS servers. Although a proprietary proxy server and firewall were in place (Microsoft Internet Security and Acceleration Server, ISA), the firewall was compromised and used as a platform for sending spam.

Once the firewall was compromised, network performance slowed dramatically and the institution's IP address was blacklisted. Uplink traffic consisted mainly of 700 Kbps of spam mail originating from the firewall. The ISA proxy server was replaced with a FreeBSD 5.4 server, configured and installed to serve as firewall and proxy (using native IPF firewall tools and Squid).

The results of this change were immediate and dramatic: Internet speed perceptibly improved as students and staff reported better performance. The network became efficient, with a maximum of only 400Kb uplink traffic. With monitoring in place, bandwidth graphs easily showed legitimate patterns and traffic. Finally, with the spam problem under control, the institution's IP address was removed from Internet blacklists.

Site Two: proxy & mail server

The institution at Site Two had 128Kbps of bandwidth, connecting approximately 50 computers networked over two sites. As with Site One, the proxy and mail server was compromised, allowing spam, peer-to-peer traffic, and viruses to penetrate. Uplink traffic was completely filled with 128Kbps of spam email.

Since the uplink was flooded, the network slowed dramatically and the institution's IP addresses were blacklisted. Again, the solution was to set up a Squid proxy on a firewalled FreeBSD server. Qmail, SpamAssassin, and ClamAV were installed on same server. ClamAV and SpamAssassin were configured to check incoming and outgoing mail viruses and spam.

As a result, viruses and spam were not transmitted, so the uplink traffic was no longer clogged. The overall Internet speed improved, and the monitoring graphs showed clean (and expected) traffic. The institution's IPs were also removed from Internet blacklists.

This server has functioned problem-free for more than 10 months and requires virtually no maintenance. Clam Antivirus updates itself every night, so no administrative interaction was required to obtain current virus signatures.

Site Three: FOSS traffic shaper

A more advanced traffic shaper was installed at this KENET site. Here, a firewall is used that is implemented with PF using ALTQ on FreeBSD. It can queue and assign priority to specific traffic types, and is implemented using FOSS.

The system is a 2 GHz Pentium IV, with 256 MB RAM, 4 NICs and 40 GB storage running FreeBSD. Even with these modest hardware requirements, the system easily handles 3 Mbps of traffic, both up and down. By using FlowC, traffic utilisation is easily broken down by application.

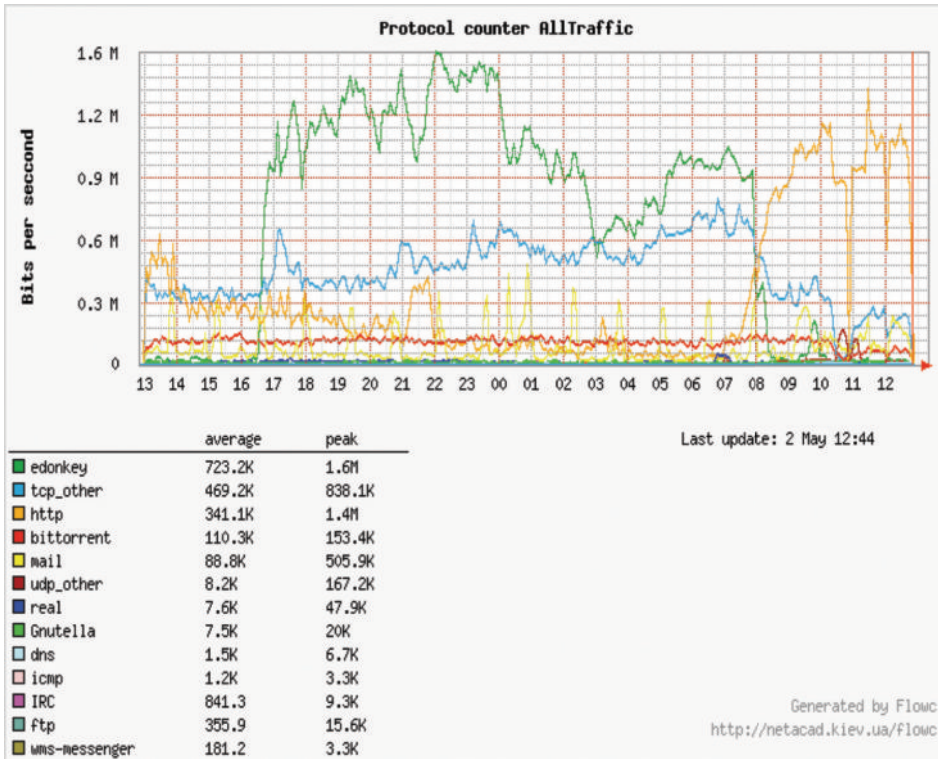

	average	peak
edonkey	723.2K	1.6M
tcp_other	469.2K	838.1K
http	341.1K	1.4M
bittorrent	110.3K	153.4K
mail	88.8K	505.9K
udp_other	8.2K	167.2K
real	7.6K	47.9K
Gnutella	7.5K	20K
dns	1.5K	6.7K
icmp	1.2K	3.3K
IRC	841.3	9.3K
ftp	355.9	15.6K
wms-messenger	181.2	3.3K

Figure 7.1: Note the dramatic shift from eDonkey to HTTP traffic at 08:00.

Notice the change in the graph when the shaper was activated at 8:00. With this FOSS server in place, efficiency in use of available bandwidth can be verified. Peer-to-peer traffic is under control, browsing is faster as HTTP now has priority, and the traffic graphs all show clean and expected traffic patterns.

--Kevin Chege

Aidworld in Accra, Ghana

In 2006, Aidworld spent three months in Ghana optimising two web portals for use over the kind of Internet connections found in research and higher education institutions there. As part of this project, we visited with members of many Ghanaian institutions to ask about their experiences using the Internet. As a result of these conversations, we gained a general understanding of the issues concerning bandwidth management and the effects of using bandwidth management strategies.

We learned that larger organisations were able to afford to employ a number of skilled network professionals to manage their network. Smaller organisations did not always have this luxury. Many institutes employ either solely part-time network administrators or none at all. The second condition is true of institutes where the network was installed by its users; people who have neither the time, nor the training, to manage such networks. Most organisations have no budget for antivirus software. Usually, Windows updates are often not installed, automatic updating is not enabled, and service packs are not applied. Existing bandwidth problems are compounded by computers infected with worms (network viruses). The computers could have avoided infection, if automatic updating and virus detection had been enabled.

As a result of bandwidth management issues, many of the institutes we visited had little or no Internet access during regular business hours. Staff members often modified their work schedules to compensate, working long, inconvenient hours in order to download important documents at night, or in the early morning.

It is often assumed that problems with Internet access and speed can only be solved by adding bandwidth. However, in our experience, we've found this is not necessarily the case. Using bandwidth management strategies are more effective than simply adding bandwidth. In the case of a network that is overloaded by virus traffic, additional bandwidth will have very little effect on the user's experience.

This case study demonstrates the need for antivirus and malware policy as a central aspect of BMO policy and the benefits of enforcing them.

The Institute for Scientific and Technological Information (INSTI) is based in Accra, the capital of Ghana. It provides information services, such as a library and Internet access, to research institutes under the Council for Scientific and Industrial Research (CSIR) throughout Ghana. INSTI acts as a hub for several organisations, providing them with access to the Internet through its own net-

work. It also provides access to visiting students and researchers through an Internet cafe. INSTI has approximately 50 staff librarians, researchers and administrators. It has about 50 computers, 20 of which comprise the Internet cafe.

At one time, the user's Internet experience of at INSTI was very slow, particularly during the day. For instance Benjamin, a librarian, would regularly have to come into work at 5:00 or 6:00 am in order to download journal articles because those were the only times that access was fast enough or available. This situation worsened until access to the Internet was stopped altogether.

Diagnosing the problem began with investigating the network traffic. It was noticed that all of the network activity lights on the switches, firewall router, and ADSL router were blinking constantly: all signs of a very high load. The traffic graphs on the ADSL router showed that the outbound bandwidth used was much higher than the connection could support, and remained so constantly. On the other hand, incoming bandwidth use was zero. That was unusual.

It was suspected that there was a problem with the ADSL router. Rebooting the router would allow Internet access, but only for a short time. The router was apparently crashing every few minutes. When unplugged from its main network, the router stopped crashing. This condition suggested an underlying cause of unusual network traffic.

The firewall (IPcop) was unable to determine which local machines were generating the traffic, or what kind of traffic it was. By inserting a laptop, configured as a transparent bridge, between the firewall and the rest of the network, it became possible to see the Internet addresses of the machines sending traffic, rather than the external address of the firewall after it had subjected them to ***network address translation*** (***NAT***). Without the internal address, it would not be possible to identify which machine was sending the traffic.

When we used a packet sniffer (tcpdump) to look at the network traffic, it was immediately apparent that most of it was UDP packets. These were sent by several local machines to a vast number of remote IP addresses, destined for one or two well-known ports. The above conditions are a classic signature of Internet worms. Most of the infected machines were in the Internet cafe. Without strictly enforcing antivirus, anti-spyware, or update policy within the Internet cafe, several of these machines had become badly infected and were flooding the network with virus traffic.

To remedy the situation, antivirus and anti-spyware software (SpyBot S&D) were installed on all machines. They were also configured for automatic Windows updates. Having made these changes, the outgoing network traffic almost reached zero, and incoming was well within the capacity of the link. Web pages loaded quickly, ping response times dramatically went down, and the router

stopped crashing. It became possible to download articles during the day which meant that the staff no longer needed to work unsocial hours in order to use the Internet.

--Alan Jackson

BMO in the UK

The following are two case studies derived from published information about Internet access in British universities. These case studies serve as useful reference points for the implementation of successful bandwidth management strategies.

JANET, UK

The Joint Academic Network (JANET) is the UK's education and research network. It connects UK universities to each other, and to the Internet. It has equivalents in many countries, such as KENET in Kenya. It serves over 18 million users, according to its website. JANET is operated and developed by a government funded organisation, the United Kingdom Education and Research Networking Association (UKERNA).

The Problem

JANET is a large network with high demand due to the number of students who gain easy, free, high speed access to it as part of their university course. They have limited funding, and in the past were not able to supply enough Internet bandwidth to meet the demand. This resulted in poor performance at peak times, and frequent outages.

The Solutions

JANET's bandwidth management strategy has two parts: an acceptable use policy (AUP), and technical measures to enforce that policy. The policy says that JANET may be used "for any purpose that is legal, that is socially acceptable to the community JANET serves, and that does not cause degradation of the performance of the network beyond that which might reasonably be expected." Degradation of network performance is exactly what bandwidth management aims to prevent.

JANET's acceptable use policy also prohibits its use for illegal file sharing, denial of service attacks, and spamming, which are some of the main bandwidth hogs on unmanaged networks. The policy gives UKERNA staff the right to shut down a JANET member who cannot or will not manage his or her own use of the bandwidth effectively, or who causes problems for JANET or its members.

A schematic map of JANET shows that it has two connections to the Internet on opposite sides of a large "ring" around the country. If that connection were to become completely full, the resulting congestion would make Internet access slow and unreliable for all JANET members. JANET has chosen to make the connections to "external links" (i.e., the Internet) equal or larger in capacity than the connections to the regional networks, which in turn serve the universities. This should mean that institutions' own connections to JANET are a bottleneck that restricts their bandwidth to the Internet and to each other, and no one institution can flood the connection.

Institutions will find that this bottleneck gives them the power and responsibility to manage the bandwidth use on their own connection. UKERNA recommends that "all JANET connected organisations should formulate a local AUP and ask staff and students to sign a declaration to confirm that they will abide by its rules" and "carry out their own internal monitoring of their network connection." A case study of how Blackburn College manages its own bandwidth over its JANET connection is given below.

JANET runs a system called Netsight. It monitors and records performance information about their network, including link status and bandwidth usage graphs. Netsight can detect and "highlight abnormal traffic levels on a site's access link that may be a result of illegal activity." UKERNA also recommends that "organisations record sufficient information about the use of their networks and maintain tools to be able to investigate and deal with problems."

JANET also offers a caching service for the Domain Name System (DNS). It is used to convert website names into addresses on the network. Caching makes DNS queries faster, more reliable and reduces the bandwidth used by DNS.

The Joint Information Systems Committee, one of the members and supporters of JANET, operates a web cache service that is used by many other members, including some that peer their own caches with it for greater efficiency. A 1996 report about this service says:

> "*The UK academic network has always offered high-performance connections within the country but comparatively slow international links. The desire to make best use of the scarce international bandwidth led to an early interest in web caching and... a national academic cache was set up... in 1995. The cache... now uses six separate computers located at sites in Canterbury and Leeds.*"

The cache has a hit rate of between 55% and 60% (documents served from the cache rather than the original website). They claim that it is one of the highest hit rates for any cache in the world, and close to the theoretical maximum. They also say that "caches act as bandwidth multipliers. For every megabyte of re-

quests arriving at the national cache, only 400 kilobytes of traffic are passed on to further networks."

In summary, JANET employs an acceptable use policy, user bandwidth limiting, and network monitoring to manage their network and bandwidth. They offer some services to their members which help the members to monitor, control, and reduce their own bandwidth use.

More information

- About JANET,
  http://www.ja.net/about/index.html
- About UKERNA,
  http://www.ja.net/about/ukerna/ukerna.html
- National Web Cache Service,
  http://www.jisc.ac.uk/index.cfm?name=acn_caching
- JANET Acceptable Use Policy (AUP),
  http://www.ja.net/services/publications/service-documentation/supportmanual/policies.html
- JANET Backbone Schematic,
  http://www.ja.net/about/topology/janetbackboneschematic.pdf
- JANET Security,
  http://www.ja.net/services/publications/service-documentation/supportmanual/security.html
- About JISC,
  http://www.jisc.ac.uk/
- JANET DNS Cache Service,
  http://www.ja.net/services/network-services/resolver/index.html

Blackburn College, UK

Blackburn College of Further Education, in north-west England, has over 10,000 full-time students. They have 1800 PCs, laptops, and printers. Their network connects to JANET for access to the Internet and to other universities.

Blackburn College incorporates disciplinary procedures to encourage users to use bandwidth wisely and in moderation. This is achieved by a combination of policy and monitoring. Technical measures to control bandwidth use are limited, but may increase in the future.

The Problem

Being a member of JANET, Blackburn College is required to comply with the Acceptable Use Policy (AUP) of JANET. This means that they have to be able to respond to complaints from JANET about abusive traffic, and track down the offender, or risk being cut off entirely from JANET.

The bandwidth use on Blackburn College's connection to JANET is approaching saturation. Due to careful monitoring they are aware of this fact before it will become a problem. Rather than upgrading their connection, they are now preparing to reduce their usage.

The Solution

As recommended by UKERNA, Blackburn College has an IT policy that covers their network and Internet connection. The college also monitors and keeps sufficient records of network usage. In the event of an investigation, they will be able to provide detailed reporting to authorities.

The IT policy includes all the terms of the JANET acceptable use policy. It makes it very clear that network traffic is monitored, that "action will be taken against users transgressing the usage codes," and that users should have no expectation of privacy.

This monitoring system gives Blackburn College the power to track down a user responsible for any problems on their network, and the IT policy gives them the ability to apply disciplinary sanctions on the user. Users are motivated to effectively manage their own bandwidth use, or else they are disconnected.

Blackburn also uses restrictive technological methods to reduce their bandwidth usage. For example, all web access from the student network and from most users on the staff network, must go through a proxy server rather than a direct connection. The proxy server denies access to specific web sites, and includes a proxy cache that can reduce bandwidth use. Blackburn filters inbound and outbound IP traffic on the border router using access control lists, on the basis of protocols and port numbers.

The College runs two caching servers, one for the staff network and one for the student network. "At the time of the case study the staff cache had a hit rate of 45% of which 25% was without any validation" (a process where the cache server makes a small request to the original web server to check that the cached copy is still valid, which takes some time and bandwidth). "The student cache had a hit rate of 40% of which 30% was without any validation." The caching servers therefore reduce their bandwidth usage by approximately 40%.

The bandwidth of the College's connection to JANET, at the time of the case study, was 4 Mbps. This was the only link on the network that was approaching saturation (maximum capacity), and risked congestion as a result. They have started a long-term collaboration with JANET's Bandwidth Management and Advisory Service (BMAS) to investigate options for reducing and managing traffic on this link.

Blackburn College will be investigating systems for monitoring the network to identify applications and protocols. They will be looking to improve their filtering, both at the gateway router (OSI layer 3) and on the proxy server (layer 7). They are continuing to investigate pre-caching of popular websites such as Microsoft.com during network off-peak times. Finally, they are investigating the effectiveness of bandwidth shaping and throttling using various products, including their existing Cisco router.

In summary, Blackburn College has a well developed and established IT policy. It is implemented using monitoring and packet filtering, which may be regarded as part of an organisational bandwidth management strategy. Their bandwidth management operations are largely reactive, manual, and policy-based, but this is expected to change as a result of their collaboration with JANET BMAS.

More information

- Blackburn College of Further Education

  http://www.blackburn.ac.uk

- Blackburn College Case Study

  http://www.ja.net/services/network-services/bmas/good-practice/part3.html

- JANET Bandwidth Management and Advisory Service

  http://www.ja.net/services/network-services/bmas/

These case studies show us that bandwidth management is necessary to get good performance from most networks. Policy and technical measures are both useful, and work best in combination, when technology is used to implement policy.

--Alan Jackson

Malawi

At the University of Malawi's Mahatma Campus in Blantyre, there are several academic and research institutions. The largest institution is the College of Medicine. It has a central campus as well as departments located in wards at a nearby hospital. In addition, there are several Malaria research facilities.

Most of these institutions are connected to the College of Medicine, using a VSAT link located there for Internet access. Previously, these institutions were interconnected by means of wireless links. The low capacity of the links, coupled with frequency disturbances (and disturbances from tress and other objects blocking the free line of sight), meant that Internet access was rather poor in most places. As a result, the 512Kb/128Kb down/up-link was rarely fully utilised. My task, together with local IT staff, was to replace the existing network, and address bandwidth management issues once the local network no longer acted as a bottleneck. We chose to replace the wireless links with fibre optic cabling.

When it came to bandwidth management, we chose to focus on three items. Firstly, we wanted to shape the Internet traffic, in part to ensure that each institution received their purchased bandwidth. Secondly, we wanted to authenticate users before they could access the Internet. At the College of Medicine, most users access the network via Windows machines connected to a Samba Domain Controller, with an LDAP database for authentication. At other sites, Active Directory was used as the domain controller. By authenticating users, we wanted to create separate bandwidth classes for staff, researchers, and students at the College of Medicine. Additionally, we hoped to prevent unauthorised users from having direct access to our bandwidth.

We decided to use tc and iptables to shape traffic, and RRDtool to graph the results. The authentication for users on Windows machines connected to the Samba Domain Controller was to be handled by a Perl script that checked which users were logged in, and opened and closed slots in the iptables firewall. For laptop users, authentication would be performed against the same LDAP database as is used by the domain controller.

How did it turn out? In setting up the new network, we chose a switched topology with VLANs for each subnet and routing done by a central routing switch. This introduced quite a few new problems. In the old network, each subnet had a gateway that filtered traffic and performed NAT. As it turned out, many machines were infected with viruses. When we removed the gateways, the viruses' rapid propagation led to congestion in the fibre network. The upstream link had an average of 100% usage, which went down to 60% when we started to filter out traffic.

In light of these problems, we decided to scrap the switched topology, opting instead for a routed network in which each subnet has its own Linux/OSPF router gateway. At each gateway, traffic is redirected to Dansguardian, before being passed to a central Squid web cache. We are currently doing some basic traffic shaping with tc and the HTB queue, while also trying to find settings that are appropriate for the VSAT link with its long delay. We are also planning additional security measures - that will automatically locate infected hosts and block them until they are clean.

To date, our case has taught us that bandwidth management requires an integrated approach. We needed to control the virus situation and prevent infected hosts from using up Internet bandwidth. We have also learned the importance of communication and policies. For example, we are performing content filtering, so someone must decide what material is appropriate. Also, there must be an efficient way for the user to alert the administrator when a useful site is blocked.

--David Blomberg, KTH

One Bellevue Center

Infospace, Inc. is a medium-sized company based in Bellevue, Washington, USA. Infospace provides the behind-the-scenes technology necessary for ring tones, mobile games, and WAP portals. To alleviate its growing pains, Infospace quickly acquired an additional lease property approximately one block from its main office. A 10 Mbps metro Ethernet service was obtained as a private Layer 2 link between the two buildings. A Netscreen 50 Firewall/VPN device was connected to the metro ethernet in each building.

While a 10 Mbps link may sound like a lot, the demand quickly outgrew the capacity engineered between the two offices. One hundred and fifty employees, each with multiple computers and VoIP telephone handsets, were scheduled to move into these offices. Security requirements included a five camera high-resolution security system, to be streamed to the physical security department located in the main building.

While VoIP conversations generally worked fine, the five security cameras consumed almost 9 Mbps of the 10 Mbps connection, and intermittently caused interruption. Internet, file sharing, and other network services came to a crawl. Network engineers were faced with a serious contention problem. They considered doubling or tripling the line rate to correct the problem. This would have increased the cost of the circuit by up to 200%. Also, it was discovered that the network provider used aggressive packet dropping (not a highly favored rate limiting method) once the data rate exceeded 10 Mbps.

However, an engineer suggested a low cost alternative: combine user education, Quality of Service, and traffic shaping while modifying several simple controls in the video security system.

First, traffic shaping was put on both Netscreens to set a maximum rate of 10 Mbps on both interfaces. This immediately solved the problem of packet drops and retries. Second, an additional traffic shaping policy provisioned 10 Mbps to VoIP services. Third, all other traffic was given a slightly lower priority, yet allowed to borrow from any of the 10 Mbps provisioned not in use. Next, network

engineers discovered that the video security system actually had video bit rate controls that could be easily changed. The system was configured in such a way that no more than 1 Mbps would be used for video streaming. Security was also informed that opening more than one camera at once would degrade network performance, and so they did not use the full 1 Mbps unless it was necessary to do so.

Users were also informed that bandwidth was limited, and to be careful with the activities they performed during business hours. This helped to make network access more fair for everyone.

All of these measures were successful in managing bandwidth within the office environment without the need for adding more capacity.

--Casey Halverson

Carnegie Mellon University

In early September of 2001, the members of the Carnegie Mellon Network Group anxiously awaited the return of the students to the campus. We knew that upon their arrival we would see a dramatic increase in our network traffic across our border router. We knew that it would be bad, indeed we were in the process of installing a new gigabit Ethernet connection to replace our OC-3, and our expectations were fulfilled. Immediately our router started dropping ATM cells. This made the situation horrible. When we dropped one 53 byte cell out of the 30 or so needed to transport a 1500 byte packet, the sending host was then required to retransmit the entire packet. This positive feedback was unbearable.

Workaround #1: Best effort rate limiting

In order to survive until our gigabit connection was ready for production, we installed committed access rate limiters on our border router. We had to experiment to find the maximum rate we could allow without dropping cells. We found that our OC-3 could only carry about 75 Mbps before it would start to drop cells. We were still dropping packets, but now they were dropping before going across the ATM link. Application latency increased dramatically, but at least the network was now usable.

Getting more than you paid for

A few months later, we had our gigabit Ethernet link up and running. We did not have to worry about oversubscribing the physical line. Our next problem was more political in nature. Carnegie Mellon and a few other Universities all contributed to the Pittsburgh Supercomputing Center's network group to manage

the "Gigapop." The PSC in turn managed a link to the Abilene high speed research network, as well as a number of commodity Internet links, via a number of tier one ISPs. Each school was then allocated a certain percentage of the overall commodity bandwidth. At the time, there were no hard limits on the maximum bandwidth. Instead, problems were dealt with personally with a phone call or email. Sure enough, it wasn't long before we heard from them, telling us that we were pushing a lot more traffic than we were paying for. It was now the spring of 2002, and we needed to come up with a new plan.

Workaround #2: Fun with rate limiting

When we installed our gigabit connection, we divided it into two Virtual LANs (VLANs) using 802.1q. One of the VLANs peered with the Internet2 / Abilene routers, and the other one peered with the commodity Internet routers. Since our bandwidth to Internet2 was not a problem, and we were required to support researchers with large bandwidth needs, we needed to constrain only traffic headed to the commodity routers. Our first thought was to use the same rate limiters we used on the ATM link. We were disappointed to find that our hardware did not support egress rate limiters. Quality of Service (QoS) would not help us, because there was no congestion on our routers.

We brainstormed a number of ideas, and our final solution was "interesting." If we could only limit traffic coming into an interface, we needed to create a point in our network where all commodity traffic would go into a single network port. Through a complex combination of VLANs, a fibre link between two ports on our border router, and redistributing Internet2 BGP routes into a new, separate OSPF process between our core and border routers, we managed to create an interface that we could rate limit.

It was not the perfect solution. We were not discriminating between different types of traffic, so every packet played Russian roulette as it attempted to pass the gauntlet of random drop. It is the kind of thing a network engineer would only do where it was an emergency and spending money was not an option. We knew that we were only buying some time to research a long term supportable solution.

More problems with packet drops

Summer came and went in 2002. During that time we dropped below our rate limiting threshold, but it was just the calm before the storm. When the students returned for the Fall session, our bandwidth usage immediately hit our thresholds and stayed there. People complained that they could not check email from home. Even SSH connections were nearly unusable. We decided to analyse the traffic usage from the previous day, and were surprised at the results. On September 4th, nine hosts were responsible for 21% of that days bandwidth

utilisation. On that same day, 47% of the traffic was easily classifiable as peer-to-peer. 18% of the traffic was HTTP. Out of the remaining traffic, we believe that 28% of it consisted of port-hopping peer-to-peer traffic. Machines on our network were acting as servers to clients outside of our campus. That was what was causing us such pain. A new solution was needed.

Requirements and considerations

We needed a long term solution, not a workaround. Here is what we needed to consider:

- The solution must not impact mission critical services (e.g. www.cmu.edu, email, ssh).
- Traffic to Internet2 must not be affected.
- We could not rely on TCP ports to punish services, since the peer-to-peer clients would hop randomly between ports.
- We needed to constrain our commodity usage to match our allowed limits.

At the time, we already knew that some P2P applications were using port hopping to get around firewalls, access-lists, and rate limiting. Indeed, we foresaw that eventually the arms race between service providers and authors of P2P software would lead to encrypting the payload and piggybacking on port 80 or 443, which could obviously never be blocked. Instead of punishing "bad traffic," we decided it would be easier to white-list "good traffic" and then make a best-effort attempt at delivery for all other packets. At first we tried to apply our rules using QoS policy, but our routers were not able to enforce them. It was apparent that any new technical solution would require a capital expense.

Researching hardware rate limiters

Our next step was to evaluate a middle-box packet rate limiter. We did a head-to-head comparison of Allot's NetEnforcer and Packeteer's Packetshaper. We found the NetEnforcer to be a better match for our requirements. We liked the method the NetEnforcer used to shape traffic: it takes advantage of TCP window size and buffering, it fairly distributes bandwidth among flows, it was easier to configure, and most importantly, it had better throughput performance when demand for bandwidth equaled the limit.

Final solution or new workaround?

In October of 2002, we put the NetEnforcer in-line with traffic and used the classification data to help develop a policy. When we returned from holiday break in January, we put the policy in place. We used a per-flow, class based,

fair bandwidth queueing policy. We had 5 classes of traffic, from high priority to low:

1. Network Critical (routing protocols)
2. Interactive (SSH and telnet) - limited per-flow, if bandwidth went over a certain limit, excess would be put into best effort queue
3. IMAP, HTTP, SMTP and other well-known ports
4. Non-classified traffic
5. P2P traffic, classified by port number

This policy resulted in fewer complaints to our help desk about accessing campus services remotely, and users' experiences seemed better. We were, however, still pushing traffic out of our network at the upper limit. We had some concerns regarding this solution:

- Per-host fair queueing was not possible. A user with 100 flows was getting 100 times the bandwidth of a user with one flow, so abuse was trivial and rampant.
- ssh -X forwarding performance was poor, making remote work difficult.
- There was high latency for UDP gaming services (which is a valid concern for a college student in the dorms).
- The demand for bandwidth was still high – socially, nothing had changed.

Application layer analysis to the rescue

In February, new software for our NetEnforcer arrived, and with it came layer 7 classification of a number of P2P services. At this time we put a hard limit on the amount of egress bandwidth these applications could consume. Our bandwidth graphs were finally showing some fluctuation below the hard limit, but the demand was still too close for comfort. We felt that we had reached the technical limit on solving our bandwidth problems. It was time for a new approach.

Social engineering

After many meetings, both in hallways and in conference rooms, we formalised our next steps. We found that we needed to change human behaviour and not network behaviour. To that end, we decided to initiate a dialogue with our campus community. Instead of the blanket email notices we had previously tried, we came up with a set of guidelines that included messages targeted to the owners of specific machines that were violating our policy. We had not tried this before because we did not have the technical systems in place to accurately

measure per host bandwidth usage. We were also unsure of how the community would react to what some might see as an unfair or unneeded restriction of "their" connection to the Internet.

But in the end, our main reason for not trying to change user behaviour is that we did not know how easy it would be.

The campus bandwidth usage guidelines

We established a daily per-host bandwidth usage limit of 1 gigabyte of data per day for traffic leaving the campus out through the commodity Internet link. As before, we did not want to limit traffic going out across our research network link. Hosts on the wireless network were given a more stringent limit of 250 Megabytes per day. We previously had hosts transferring multiple gigabytes of data in a single day over their wireless link. We published these guidelines and requested comments from the campus community. We also included an exemption for specific hosts. The exception clause in our policy states:

> "*Owners of servers or services who exceed the usage guideline and feel they have legitimate cause to do so, must contact Computing Services. In rare cases an exemption may be granted. In all cases, solutions for fair use of the Commodity Link bandwidth will favor those services that support the research and educational goals of the university.*"

Our goal was to help the users follow the rules. To do this, we came up with a system for enforcing them that gave the users a chance to change their ways. We wanted the overall experience of the users to be positive. To accomplish this, our enforcement system worked this way:

- If a host exported more than 10 gigabytes of data over a 5 day period, we sent out an initial email message that informed the owner of the quota and how much they actually used. We requested that they provide information regarding their usage, if they felt it was legitimate, otherwise they should decrease their bandwidth consumption. If they needed help, we gave them our help desk's contact information.
- We then waited two days before taking any further action. On the third day we would start checking their daily usage and comparing it to the quota. If they ever exported 2 gigabytes or more in a single day, we would send a warning message. If the initial notification was the "good cop" message, this one played the role of "bad cop." The email warned the user that they must either provide a reason why the host should be exempted from the guidelines or decrease their bandwidth usage. If they did not comply, we would disconnect their host from the network. This would be treated like any other network abuse issue.

- We would then wait two more days to give them a chance to change their usage behaviour. After that, if they exceeded more than 2 gigabytes in a single day without contacting us to explain the situation, we would carry out our disconnection.

Human effort

Initially, the system was carried out manually. We had an automated system that would generate a daily report of all hosts and how much bandwidth they used, sorted by top usage. We already had a system in place that could tie every MAC address with a specific user or group. Then it was just a matter of going through the list by hand, updating a database of hosts in the various warning states, and sending the email messages. This took about two hours per day for one engineer. In the first week, we sent out 49 initial inquiry messages and 11 warnings. In March, we sent out 155 initial messages, 26 warning messages, and disconnected 5 hosts.

You may have have noticed that we were acting on numbers that were twice that of the policy. We thought it would be easier to find the sweet spot of per host limits if we started enforcing at a higher level and worked down to the official quota. Indeed, it was good that we did. On the first day of enforcement, we dropped our Internet usage from 120 Mbps to 90 Mbps. We were no longer dropping packets with our hard limits, and we were convinced that we had found a solution that worked. Our next step was to spend the resources to develop a software solution that would automate the entire system.

Positive results

We received very few negative responses related to the guidelines. Most users were completely unaware of their bandwidth usage. Many had installed peer-to-peer applications that ran in the background, unbeknownst to them. Some users discovered trojan FTP servers running on their hosts. Some of my favorite email response were "I don't understand--what's bandwidth?? How do I reduce it? What am I doing wrong? What's going on?!" and "I have no idea what bandwidth is or how one goes about using it too much."

By the summer of that year, we sent a total of 489 initial inquiry messages and 102 warnings. We revoked network access to 29 hosts and exempted 27. Over 39 machines were found to be compromised.

Conclusion

This system worked very well for us. There were two systems we had in place before we started working on the problem that proved to be invaluable. One of them was a system that allowed users to register their hosts on the network.

This self-service system provided the database that tied a MAC address to a specific user. It is called NetReg (available at *http://www.net.cmu.edu/netreg/*). Another system we had in place was an Argus monitoring infrastructure (*http://www.qosient.com/argus/*). This is a set of hosts that collects network flow data from routers. This was used to generate the "top talker" network usage reports. Both of these packages are freely available from their developers. Both of them require a decent system administrator to get up and running.

The big lesson we learned was that users were willing to change their behaviour, and that without their cooperation, technical solutions alone could not solve the problem. It was worth the effort to work with them to manage the bandwidth at our organisation. If you would like more information, here are some interesting links:

- Carnegie Mellon Network Bandwidth Usage Guideline: Wired network *http://www.cmu.edu/computing/documentation/policies_bandwidth/bandwidth.html*
- Carnegie Mellon Network Bandwidth Usage Guideline: Wireless network *http://www.cmu.edu/computing/documentation/policies_wirelessbw/wireless_bandwidth.html*
- How Computing Services Monitors and Enforces Network Bandwidth Usage *http://www.cmu.edu/computing/documentation/bandwidth_howenforced/index.html*
- Internet2 Joint Techs presentation resource page: "Successful Bandwidth Management at Carnegie Mellon," *http://www.net.cmu.edu/pres/jt0803/*

--Peter John Hill