

Review of Bandwidth Management Technologies, Availability, and Relevance to UK Education

Version 1

December 2003

Dr J. P. Knight

Loughborough University Computing Services/

JANET Bandwidth Management Advisory Service

1. Introduction.....	7
2. Bandwidth Management Concepts	9
3. Caching	10
3.1 Introduction.....	10
3.2 Deployment Considerations.....	10
3.3 Peering and Hierarchical Caching	11
4. Rate Limiting, Packet Shaping and Quality of Service	13
4.1 Rate Limiting	13
4.2 Shaping	13
4.3 Common Algorithms	14
4.3.1 First In, First Out (FIFO)	14
4.3.2 Random Early Detection (RED)	14
4.3.3 Token Bucket Filtering (TBF)	14
4.3.4 Stochastic Fairness Queuing (SFQ)	14
4.3.5 Class Based Queuing (CBQ)	14
4.3.6 Hierarchical Token Bucket (HTB).....	15
4.3.7 Alternatives	15
4.4 Quality of Service (QoS)	15
5 Compression	18
6 Content Filtering	20
7 Access Control Lists	22
8 Firewalls.....	24
9 Good Practice/Netiquette	26
10. Available Bandwidth Management Technologies and Implementations	29
10.1 Caching	29
10.2 Web Servers	31
10.3 Rate Limiting and Packet Shaping.....	32
10.4 Quality of Service	33
10.5 Compression	34
10.6 Filtering.....	34
10.7 Access Control	35
11. Costs, Popularity and Relevance of Products to FE/HE	37
11.1 Caching	37
11.2 Rate Limiting and Packet Shaping.....	37
11.3 Quality of Service	37
11.4 Compression	37
11.5 Content Filtering	38
11.6 Access Control Lists	38
12. Future Directions	39

JANET Technical Guides

JANET Technical Guides are one of a series of user guides available to JANET customers. Technical Guides contain detailed technical information and are intended for technical support staff at JANET sites, network specialists or those with a particular interest in the specialist area.

If you have any queries or comments about the Technical Guides or would like to obtain copies, please contact:

JANET Customer Service

UKERNA

Atlas Centre, Chilton, Didcot

Oxfordshire, OX11 0QS

Tel: +44 (0) 1235 822 212

Fax: +44 (0) 1235 822 397

E-mail: service@janet.ac.uk

Further details of the documents in this series are available at:

<http://www.ja.net/documents/>

1. Introduction

Whilst in general prices for network bandwidth are falling over time, the demand for bandwidth within organisations is constantly rising. This demand is caused by the increased use of electronic resources for teaching and learning, the spread of remote research collaboration and the spread of common desktop applications that will use practically any amount of bandwidth given to them. Not all of the last category could be described as having much academic worth and indeed some may be viewed as undesirable by many institutions (e.g. games, file-sharing, etc.).

It is now recognised that one of the tasks that network teams within Further Education (FE) and Higher Education (HE) need to tackle is the management of bandwidth. There are many techniques available to them and many competing demands for the bandwidth available. This report aims to provide an overview of these techniques, look at those products and technologies available and aid institutions in choosing between them. The associated costs, popularity and relevance to FE and HE are also detailed. Finally, a likely set of short term future directions and longer term possibilities is provided.

This report is from the JANET Bandwidth Management Advisory Service - BMAS. For more information about BMAS and updated versions of this document, please visit our website:

<http://bmas.ja.net/>

For more information about education and research networking in the United Kingdom, please see the Joint Academic Network (JANET) website:

<http://www.ja.net/>

2. Bandwidth Management Concepts

Bandwidth management is not a single technique or tool. Successful provision of managed network bandwidth within an organisation is likely to involve the application of many tools encompassing a number of different techniques. The techniques and tools an institution uses will depend on a number of factors:

- ratio of available bandwidth to existing/future demand;
- need to prioritise some traffic types/users over others;
- resources available to implement bandwidth management strategies;
- organisational experience with products and systems.

Before an institution can decide how to address the problem of reducing bandwidth demands, it is important to be aware of how bandwidth is being consumed at the moment. There are a number of options for doing this. Firstly some of the bandwidth management products mentioned in the next section have a monitoring as well as proactive mode (for example Packeteer has the PacketSeeker monitor as well as the PacketShaper control system). Unfortunately these products are often expensive and are rarely available for loan (as sites could just monitor their traffic for two weeks and then hand them back!).

Luckily packet sniffers that run on normal hosts are widely available. Many network managers will already be familiar with products such as Ethereal, Windows Network Monitor and tcpdump. Packet sniffers are often cheap or even free and just require a PC with a good network card suitably placed on the network. Some network switch and router vendors also have offerings that are included in their products that can help (such as port mirroring and NetFlow). Simple Network Management Protocol (SNMP) based network management solutions will also give an overall indication of bandwidth 'hotspots' on the network that can be further investigated.

3. Caching

3.1 Introduction

Caching resources allows a single copy of a resource to be downloaded over an external network connection and then served out to multiple users locally. Caching not only reduces the amount of bandwidth used on the external network connection but can also sometimes provide increased performance for the local users. After the initial download of the resource the users do not have to compete for bandwidth on what is likely to be a slow and congested external network path. Instead, users will get a copy of the resource delivered rapidly to them from a local cache server over what will normally be a much faster and less heavily loaded internal institutional network connection.

Caching techniques have been deployed for many protocols in use on the Internet, in particular the Hypertext Transfer Protocol (HTTP) and File Transfer Protocol (FTP). Web caches are servers that run a proxy cache application that takes the requests for objects (web pages, sounds, images, movies, etc.) from web browsers on client machines and checks for them in a large file-store of previously retrieved objects. If the required object is present and up-to-date in the cache file store it is returned straight to the web browser, otherwise a request is made to the remote server hosting that object and it is cached as it passes through the proxy cache on the way to the client browser.

3.2 Deployment Considerations

There are a number of issues to bear in mind with web caching. Firstly, not all requests and/or objects are cacheable. For example Secure Socket Layer (SSL) encapsulated HTTP sessions (https:// URLs) are normally not cacheable. They often include personal data and the end user can receive error messages indicating that the connection to the server has failed. Audio and video streams also require special handling due to real time constraints. HyperText Markup Language (HTML) form submissions and the resulting dynamically generated responses often cannot be cached.

Many commercial database providers use Internet Protocol (IP) based authentication techniques. Whilst many would argue that this is an inherently broken means of authentication (it being relatively easy to spoof IP addresses) it is in widespread use. Web caches cause problems in a number of ways. Firstly they might be located within the 'valid' range of IP addresses accepted for a remote service and yet they themselves might service a much larger pool of IP addresses. This would result in users that the vendor was not anticipating gaining access via the web cache. If an organisation is billed per access this could mean running up hefty bills providing an unexpected service. On the other hand, some IP address authenticated services specifically look for the signs of proxy caching in use and refuse to service the request; in these cases the only way to access the remote service is to force the requests to bypass the proxy cache.

Web caches also have to be considered in terms of a site's overall security policy. Open web caches that will accept requests from all comers are in demand from malicious Internet users who can use the server to hide their identities whilst attacking remote web sites or laundering connections to other protocols through them. Web caches can also use proxy authentication to check user identities and ensure that only the members of an institution can access the proxy from off site. This has the added bonus of giving their web traffic IP addresses in the range assigned to the organisation. This can be useful if the institution subscribes to many databases with IP address based access controls.

3.3 Peering and Hierarchical Caching

Caches can also be used in hierarchical schemes. In a hierarchical scheme local web caches can talk to 'siblings' (other web caches that are the same level in the hierarchy as they are) and 'parents' (web caches nearer to the network core that aggregate caching requests from many child caches and thus huge numbers of end users). The 'parents' themselves can be children of other caches higher up in the scheme. Hierarchical caching has the advantage of being able to attain a relatively high hit rate due to the large user base. However it needs to be scaled carefully or else users will experience increasing delays resulting from requests and responses traversing many levels of caches. Techniques have been developed to allow caches to provide each other with summaries of the objects they hold in an effort to short circuit these hierarchy traversals.

Of course to make a useful impact on bandwidth usage within an organisation the use of caches has to be strongly encouraged and possibly even enforced. Many institutions have taken to supplying web browsers pre-configured to use their campus proxy cache and/or making use of features available in some routers to transparently redirect web traffic through the cache. However it should be noted that proxy authentication can only be used if the proxy cache is configured explicitly in the browser; it does not work if proxy interception (or 'transparent caching') as it is sometimes known) is in use.

During periods of lower bandwidth utilisation, web content pre-fetching can be used to refresh cached objects and retrieve perceived content needed for teaching and research. Pre-fetching web content can reduce the perceived latency for the retrieval of some web sites and reduce the bandwidth requirements during peak usage. When configuring pre-fetching, it is necessary to restrict the time window to when the link is using relatively little bandwidth (for example during the middle of the night). Pre-fetching content into caches obviously relies on the content being fairly static. Dynamic web pages and constantly updated websites (such as news portals and weblogs) will not provide much benefit from pre-fetching into the cache.

4. Rate Limiting, Packet Shaping and Quality of Service

4.1 Rate Limiting

Sometimes the amount of traffic wanting to use a network link far exceeds the available capacity of the link. Rate limiting techniques are designed to force some protocols to reduce their demands for bandwidth based on the protocol, network interface or user involved. Traffic shaping on the other hand aims to spread out the demand for bandwidth and ensure that the most efficient use is made of the available bandwidth. Both of these help to either reduce the required bandwidth to within the available capacity or to give competing traffic a chance of using the link.

Rate limiting can be applied on a per-protocol, per-interface or even per-user basis. It can take place both in the network (in switches and routers) and also within servers (e.g. throttling the bandwidth served out by web servers). Traditionally network based bandwidth rate limiting was applied to a particular port (e.g. port 25 for Simple Mail Transfer Protocol (SMTP), port 80 for HTTP, etc.) or across particular subnets. This worked well enough until applications started to share the same protocols and/or ports. For example many applications beside traditional web servers now use HTTP and port 80, such as peer-to-peer filesharing systems. Institutions may wish to prioritise traditional web traffic over these other systems.

4.2 Shaping

Modern traffic or packet shapers on the other hand delve into the IP packets and attempt to deduce information about the application layer source or destination of the traffic. These traffic/packet shapers effectively use a form of pattern matching on the Layer 7 data enclosed in the packet in order to assign the packet to a particular 'flow' that then uses a given shaping rule set. To distinguish between peer-to-peer traffic and normal web traffic one needs to examine the contents of the packets and find a pattern to match that will disambiguate. It may be necessary to re-assemble fragmented IP packets into flows in order to find this information, since the pattern being matched can be split over two or more fragments.

Once the packet classifier has determined which application and thus which rule set to use, there are a number of different ways to shape the traffic into the available bandwidth. The rule sets might provide a fixed breakdown of available bandwidth between applications so that, for example, a real web HTTP session on port 80 gets 50% of the available bandwidth, a real SMTP or Internet Message Access Protocol (IMAP) e-mail session gets 25% and the remaining 25% is contended for by all other applications, even if they are using ports 25, 80 or 143. As another example, the spread of Voice over IP may encourage network administrators to dedicate a known proportion of their bandwidth to trunk voice calls between telephone switches.

Alternatively, varying degrees of dynamic adjustment may be employed to allow non-preferred applications to take up bandwidth when preferred applications are not using it. Thus unwanted applications may be throttled down to virtually nothing, effectively blocking them.

4.3 Common Algorithms

Many rate limiting techniques operate on queues of packets formed from the classified flows. There are a large (and growing!) number of queue management techniques:

4.3.1 First In, First Out (FIFO)

This is a simple strategy that simply queues up packets when congestion or bandwidth limits are hit. The earliest packets to arrive will be the earliest ones to leave. FIFO queues do not have any internal ways of altering the priority of packets within them and if they fill up they will suffer from 'tail drops' where the most recently received packets are discarded even though older packets of a flow are still in the queue).

4.3.2 Random Early Detection (RED)

These algorithms work by randomly throwing away packets when congestion starts to appear or bandwidth limits are approaching. These work best with application protocols using underlying reliable transport protocols such as Transport Control Protocol (TCP). This is because the source will not only notice the lost packets and retransmit them later, but will also reduce the window size to slow the whole flow down. Some RED variants (of which there are many) introduce weightings so that the probability of discarding packets varies depending on which queue a flow gets put into - low priority/over subscribed traffic flows are put into queues where packets are more likely to be discarded and thus will be the first to be shaped.

4.3.3 Token Bucket Filtering (TBF)

A TBF algorithm prevents high-priority traffic from totally starving lower priority flows of bandwidth. It does this by only passing packets arriving at a rate that does not exceed a specified level set by the network manager. It is very precise and has a low demand on Central Processing Unit (CPU) cycles so it is useful on low end hardware, but does not offer the flexibility of some of the other queuing mechanisms.

4.3.4 Stochastic Fairness Queuing (SFQ)

This distributes flows into a set of FIFOs using a hashing algorithm and then visits each FIFO in turn.

4.3.5 Class Based Queuing (CBQ)

This is designed to reduce the output to the required rate by enforcing idle periods in the output queue by calculating what the average time between packet arrives should be to give the desired output rate. It then subtracts this from the Exponential Weighted Moving Average (EWMA) idle time of the real link. The EWMA treats new packet arrival times as exponentially more important than previous ones so that recent traffic has more effect than old traffic, but the figure is still based on everything that has been seen. The result of the subtraction is then used to do the queuing. A zero result shows a perfectly balanced link, a positive value indicates an under utilised link, and so no idle time needs to be injected, and a negative value indicates an overloaded link that needs throttling.

4.3.6 Hierarchical Token Bucket (HTB)

This is a classful subset of CBQ that has less flexibility, but is simpler to configure and requires fewer systems resources to support. HTB ensures that a class of traffic is allocated the minimum amount of bandwidth it asked for. As the name implies, HTB consists of a collection of TBFs placed in a tree-like hierarchy. Each TBF controls one queue that has one or more flows allocated to it.

4.3.7 Alternatives

Some rate limiting tools can also make use of protocols' inherent bandwidth management options. For example one might attempt to throttle the bandwidth used by TCP based streams by persuading a source that its window needs to be reduced. This can be achieved by sending the source of a flow a carefully crafted TCP packet without the shaper ever doing any queuing of the flow's traffic.

4.4 Quality of Service (QoS)

Sometimes a network manager may wish to offer different levels of service for different subsets of the traffic passing over the infrastructure. QoS technologies allow the manager to do this. The service levels that QoS deals with include dedicated bandwidth, reduced packet loss, controllable amounts of jitter and latency and ensure that particular traffic flows do not swamp the network and drown out other flows. These facilities can help ensure that critical traffic is delivered more reliably and that real time traffic such as video and audio conferencing work more effectively.

To make use of QoS, packets in flow need to be identified, classified and then marked. The identification and classification of packets is performed as described above for traffic shaping, either by source, destination and ports or by using application level probing into the packet contents. However unlike simple traffic shapers that just alter the flow's characteristics at a particular hop (mostly by delaying or dropping certain packets whilst letting others through first), the QoS classified packets are marked in some way to allow multiple hops through the network the option of altering the way they are handled without having to identify and classify them repeatedly.

Once packets have been classified and QoS marked, the end-to-end QoS tools embedded in the network can then provide several levels of service to them:

- Best Effort. This is the same as having no QoS at all and is what most networks provide by default.
- Differentiated. This is sometimes called 'soft' QoS or DiffServ. Some traffic gets some form of preferential treatment over the rest, although it does not offer any guarantees. This is a statistical bias towards certain types or flows of traffic when queued in switches or routers.
- Guaranteed. This is sometimes called 'hard' QoS or IntServ. It reserves resources in the network for particular flows. Those resources might be chunks of bandwidth, paths through a switch fabric, short cutting of queues, etc. Once a flow has an end-to-end hard QoS set up for it, it will definitely get its required performance from the network under normal conditions.

The tools used to implement these service levels provide for congestion avoidance and management, as well as link use efficiency to manage delay on low speed or congested links. QoS implementations can use the bandwidth throttling and traffic shaping technologies mentioned above at each node (host, router, etc.). The important point is that each node in the flow is co-operating rather than acting alone.

Differentiated QoS uses a precedence marked on the IP packets using the three precedence bits held in the Internet Protocol packet header – the Differentiated Services Code Point (DSCP), originally known as the Type of Service (ToS) field. These three bits can specify six classes of traffic plus two additional reserved patterns. DSCP values are normally set and used by network devices such as switches and routers to decide the type of service that will be available to the flow in question.

The Internet Standard Resource ReSerVation Protocol (RSVP) provides a set up mechanism for a host to request a specific QoS be applied to an applications data flow through the network. It has also been implemented within some networking devices in such a way as to allow network managers to use it to reserve resources for non-RSVP aware hosts/applications. It can be used to set up resources for hard QoS systems. The QoS and policy information that RSVP carries through the network is opaque to the protocol itself. It is used by admission control and policy mechanisms within the network to determine how to handle the application data flow.

RSVP operates as a transport level protocol on top of Internet Protocol version 4 (IPv4) or Internet Protocol version 6 or (IPv6). However rather than transporting application data, it transports messages to reserve resources along the route of a flow. RSVP is not a routing protocol itself; the routing protocol(s) in use determine where each packet of a flow goes whilst RSVP determines what QoS those packets should be marked with when forwarded.

RSVP is designed to handle unidirectional flows; to reserve resources in both directions (for example a two way Internet telephone conversation) would require two RSVP instances, one for each direction. It is also intended to be used from the receiver of a data stream to the source, rather than from the source outwards. This may seem strange until one considers that it is designed to work with multicast groups where there may be many receivers for a single source. RSVP requests from the receivers will propagate along the reverse path to the application data until they reach the router where the receiver's individual data flow is broken off from the large multicast tree.

5 Compression

Data compression is an obvious way of reducing the bandwidth required for many traffic sources. Some servers provide the option of allowing arbitrary data streams to be compressed on the fly when talking to clients that can handle the compressed data streams. This allows objects to reside on the server in uncompressed, easily manipulated forms whilst still saving bandwidth. Some vendors also support compression in hardware, e.g. to accelerate an HTTP server or a slow leased line network link.

Some protocols permit their headers to be compressed on the fly. For example the Point to Point Protocol (PPP) used for dial up services often has header compression turned on so that it reduces the bandwidth demands it makes on slow telephone connections. The Real Time Protocol (RTP) can also be compressed in flight which can be a real boon given the time critical nature of the flows it forms (typically multi media conferencing type applications). The RTP protocol is layered on top of User Datagram Protocol (UDP) and IP resulting in a 40 byte header that is then attached to a data payload of between 20 and 150 bytes. In the worst case the packet is 300% of the data! RTP compression reduces the header to around two to five bytes, and thus achieves quite a considerable saving in bandwidth.

For data payloads one has to consider whether the compression algorithm used is lossless or lossy. A lossless algorithm can take a data stream, compress it down and then uncompress it later back to an identical copy of the original data stream. Lossless compression is used to compress program code, documents and medical images where the original data stream must be recovered intact. Examples of such compression mechanisms include GZIPed files, the many PC and Mac based file compression and archiving tools, and Graphic Interchange Format (GIF) image compression.

Lossy compression schemes do not produce the same data stream from the decompression system as was compressed originally. These algorithms are often employed for audio, video and images, where the human brain can often still make sense (or not even notice) the degraded quality of the output. Typically they offer a larger amount of compression traded off against the reduced output quality. Some formats allow the amount of compression or degradation to be varied by the user at compression time. Examples of these types of compression algorithms include Joint Photographic Expert Group (JPEG) encoded images, and Moving Picture Expert Group (MPEG) audio and video streams.

Some data formats can support multiple compression mechanisms to allow the user to decide upon the appropriate compression mechanism for a particular application. For example the Silicon Graphics movie file format can carry either raw uncompressed video frames for situations where bandwidth is not as important as preserving the original video frames, or one of two different compression mechanisms giving differing bandwidth/quality trade offs.

Lastly it should be noted that if a data stream is already well compressed, attempting to compress it again will often result in a larger output stream. This is because compression algorithms are looking to replace repeating patterns in the data stream with smaller tokens to represent them. The resulting compressed data has few if any large patterns in it (often appearing to be nearly random data) and so the compression algorithm overheads just make a second compression produce a larger file. Many static file compression algorithms implementations will spot this and discard the second compression. That is more difficult to do when data is being compressed on the fly, unless the server doing the compression detects the existing compression algorithm's header and then passes the pre-compressed data through unchanged.

6 Content Filtering

Many sites make use of content filters. This can be to protect minors from obscene or other 'harmful' material. In other situations filters are used for a less headline grabbing role to cut out transfers of data to/from sites that the institution's administrators have determined offer little or no resources of worth to the organisation's stated aims. The bandwidth that would otherwise be wasted can be used for traffic that does further those aims. Lastly, filters can also be used to cut out really harmful data such as viruses and Trojan Horse programs.

Content filters can operate on a number of different levels. At the lowest level they are all based on pattern matching of some sort. This might be as simple as a web proxy cache matching a requested URL against a list of banned sites and can get as sophisticated as filters that determine whether images contain sufficient flesh tones to make them likely to be pornographic. Some filtering systems make use of scoring systems where multiple patterns are matched in a resource and each pattern is given a different score or weight. The total score for the resource then determines how it is treated.

Web proxy or cache server filters are probably some of the most widely implemented and controversial filters available. These are used to determine whether requested web traffic should be retrieved for the end user and passed through intact, if it should be retrieved but have some patterns, typically words or phrases removed or replaced, or if the request should either be denied or redirected to another URL. Filters typically make use of one or more of a number of heuristics:

- pass lists (resources that are explicitly allowed);
- block lists (resources that are explicitly banned);
- scoring on text patterns (so that swear words for example might have relatively high scores but not high enough for an individual word to block a page);
- scoring based on analysis of images or hyperlinks.

Some commercial products are effectively 'black boxes' that are inserted into the network with predefined pass and block lists incorporated into system. In this instance it is strongly recommended that network administrators assess how much editorial control is available to them for the maintenance and upgrading of such lists.

As well as web filters, mail system filters are increasingly popular. These can be used to cut down the levels of Unsolicited Bulk/Commercial E-mails (UBEs, UCEs) or 'spam' as these messages are commonly known. Again pattern matching and scoring is used, coupled with widely distributed lists of known spam hosting sites. Some systems make use of adaptive Bayesian filters that can be trained to spot spam and other unwanted e-mail by providing the filters with a large pool of pre-classified e-mail. Currently such filters have hit rates of over 90% and relatively low rates of false positives (where a legitimate e-mail

is filtered as spam by mistake). However there is ongoing competition between the spammers and the writers of anti-spam filters.

As soon as an anti-spam filtering system becomes popular it tends to be attacked by spammers using a combination of litigation (usually on the grounds of free speech) and technical work to defeat the technology. Denial-of-service attacks (discussed in detail below) are also becoming increasingly popular. Once circumvented, the increased flow of spam causes the filter writers to create a better set of filters and the cycle starts again. Some products provide off-site filtering for institutional e-mail so network administrators do not have to keep up. E-mail filtering can also help in the battle against viruses. There are now several antivirus products that can scan both incoming and outgoing e-mails. Scanning incoming e-mails helps to stop clients on the site from being infected, which might then result in a large consumption of bandwidth internally on the site as they send viral payloads to anyone in any address book they can find. The outgoing e-mail filters can stop internal infections from leaving the site and thus consuming external link bandwidth.

Of course e-mail filtering can also be passed on to the end users if required, to allow them to manage the increasing influx of e-mail that most of us are now experiencing. This is not a bandwidth saving as such but it can be a popular side effect of implementing a bandwidth saving filtering system.

7 Access Control Lists

Access controls allow an administrator to determine which users or traffic flows are permitted to use which resources. Access Control Lists (ACLs) are familiar to many multi-user systems administrators and similar options are available in most web servers, web caches, general purpose file servers, e-mail systems, switches and routers.

In a web server, ACLs can be used to limit the ability of users to retrieve certain resources. These users may be identified by a user name and password combination supplied to the web server by their browser, or the subnet or IP address that the traffic is originating from or even the method or URL that they are using to access the resource. The ACLs can be used as a simple allow or deny security mechanism or can be combined with the web server's rate limiting or on the fly compression facilities to provide access to a resource but with different bandwidth and delay characteristics for different users.

Web managers should also consider the use of the `robots.txt` file the de facto standard for specifying where the spidering robot is allowed to go on a web site to limit where well-behaved spidering robots will go. The spidering robots are used by indexing services such as Google and Altavista to retrieve pages by following the chain of hyperlinks. Spidering robots that ignore the `robots.txt` are not well behaved and should be considered for banning, especially on a large web server. Not only can they use up bandwidth unnecessarily, they will also make use of information that may be out of date or inaccurate and can also be trying to harvest e-mail addresses for spammers.

Routers and switches typically use ACLs at a lower level to control features such as Network Address Translation (NAT) and Port Address Translation (PAT), membership of Virtual Local Area Networks (VLANs), Dynamic Host Configuration Protocol (DHCP) pools, etc. and identify flows that can be passed on to QoS and traffic shaping facilities. They can also be used to limit access to the management interfaces of switches and routers to reduce attack vectors to which some devices are susceptible.

8 Firewalls

In a way, a firewall can also be viewed as an access controlled content filter. It is designed to look at the source, destination and/or content of flows and determine, which flows to allow through and which to drop and log based on ACLs. Firewalls are probably one of the most common bandwidth management tools deployed today, with many academic sites using them as well as companies and individuals. They can be a useful first line of defence against attack as well as helping to reduce bandwidth requirements. However it is worth bearing in mind that institutional firewalls will sometimes stop legitimate academic traffic and so network managers must be prepared to liaise with end users over the filter rules that a firewall implements.

When configuring firewalls it is worth remembering that the duplex flow in normal network communication can have traffic take different paths through the network on the outward and return journeys known as 'asymmetric routing'. Firewall rules need to be carefully constructed to take account of that.

Some ACL systems include the concept of time in the ACLs as well. This can be used to alter the filter rules for different times of day or to permit different traffic patterns at weekends. For example a site might block the use of web-based e-mail services in teaching labs during the day when students could get distracted by personal e-mail, but open it up at night and weekends.

The firewall is typically thought of as a device that sits in between the campus network and the Internet and mediates access. In practice it is possible that multiple firewalls or router ACLs may be in place between a host on the campus network and another host elsewhere on the Internet. Many institutions and Internet backbone carriers apply their own routing policies, e.g. to mitigate against spam and worm/virus propagation.

Most modern desktop and server operating systems also provide their own firewall subsystem. Some of these are crude and only provide very basic protection, such as the Windows XP built-in personal firewall. Others are industrial strength and actually used in a number of commercial products, such as the Linux NetFilter (iptables) subsystem. Host based firewall subsystems can be used to 'hide' hosts or services running on them from parts of the campus network - or from the Internet at large – even when no formal firewall or router ACL set exists for an institutional network.

9 Good Practice/Netiquette

There are several techniques that the careful network manager and/or end user can implement to help with reducing the bandwidth demands on campus. A brief overview of some of these is given here.

Multicast protocols should be used for widely viewed audio-visual streaming media rather than multiple unicast streams. A multicast stream is only split in the multicast aware network topology when it actually needs to go two different ways. This means that a single stream can leave a site and be split three ways within the Regional Network that the site is connected to. Two of those streams could then enter two other organisations that are connected to the same Regional Network and be split again internally to serve a number of users spread over each site. The third stream within the Regional Network can then be passed up to the JANET core for distribution to other Regional Networks or further afield outside of the JANET community. In this scenario each organisation sends and receives just one multicast stream. If the same multimedia stream had been sent with multiple unicast streams, each receiving site would have had the same set of bytes arrive once for each user. This could result in the transmitting site having a huge number of near identical outgoing streams. When one considers the transmission of a popular workshop or conference to a large number of geographically dispersed viewers, the use of multicast can be seen as a big bonus in terms of bandwidth utilisation.

Web managers at institutions can also play their part in helping reduce bandwidth demands and at the same time improve the usability of their web site. The main culprit is the over-emphasis on images and other embedded multimedia content. Pages that use multiple graphics should be kept to a minimum, especially if the page is likely to be viewed over a slow modem links. This is especially important if the web site is being used to sell something to home users and schools; not everyone has broadband yet. Text-only pages should be available where possible. This not only reduces bandwidth demands but can also be easier for some users to navigate. In the UK the Disability Discrimination Act has made it mandatory to design web sites that are accessible to users with disabilities, e.g. a visual impairment.

Sites that have a Usenet News feed may wish to check that the newsgroups they are receiving are genuinely of interest. Usenet is responsible for a relatively low percentage of traffic these days but can still contribute several hundred megabytes of traffic a day. It may also be desirable to filter Usenet traffic for the same reasons that e-mail is filtered.

If a host on an institution's network falls victim to a virus, worm or Trojan Horse program, or is compromised, the network manager should quickly take steps to ensure that the machine is disconnected from the network (going as far as disabling the switch port that the machine is connected to if the user is either unavailable/unwilling or unable to disconnect the machine themselves). This not only prevents further damage being done to that machine and other

ones within the institution's Local Area Network (LAN), but also prevents the machine launching potentially bandwidth intensive attacks against other sites. These may be as a result of the virus trying to spread or it may be that a compromised machine is either trying to crack further machines at other sites or is involved in a Distributed Denial of Service (DDoS) attack against another site.

These attacks are potentially very serious. They can quickly use all the bandwidth to a site, cause routers and hosts to run out of CPU cycles and generally disrupt other network protocols. Network managers who find themselves the target of a DDoS attack (rather than just an unwilling participant) will have a major problem and will need to liaise with upstream service providers such as their Regional Network. With any security related incident it is also good practice to involve the JANET CERT:

<http://www.ja.net/cert/>

An institutional firewall will provide routine protection for the campus network, and can help prevent some of the worst effects of a DDoS attack. Most systems will also provide logging information that will assist in tracking the sources of the attack – although it is common for attackers to forge IP addresses. A networked intrusion detection system can also give early warnings of impending attacks - a number of proprietary and open source intrusion detection systems are available.

Firewalling and bandwidth management services have often been provided using software on commodity PC based server hardware and operating systems. This approach has limitations imposed by the operating system design and hardware restrictions such as bus bandwidth. Hardware assisted products using Application Specific Integrated Circuits (ASICs) are available, but these have yet to become commodity items and are significantly more expensive than software based solutions.

Network performance on commodity server hardware may be enhanced using TCP/IP Offload Engine (TOE) adaptors, which use custom ASICs to carry out many common networking tasks that would otherwise occupy the CPU of the host server. However, note that the TOE products available at the time of writing are not suitable for use in firewalls or other applications where the host server carries out arbitrary transformations on received packets. A typical use of TOE would be to accelerate a busy web or e-mail server, or web cache.

10. Available Bandwidth Management Technologies and Implementations

What follows is an introduction to some of the products that implement the techniques described in the previous section. This list should not be viewed as definitive; new products are appearing daily and existing products reach their End Of Life. Please note that the JANET Bandwidth Management Advisory Service does not endorse or directly support any products - institutions are recommended to evaluate the competing products themselves in light of their own requirements.

10.1 Caching

Apache mod_proxy:

http://httpd.apache.org/docs/mod/mod_proxy.html

A module for the world's most popular web server that provides proxying capabilities for HTTP, FTP and SSL CONNECTs. It can act as a web cache, can pass requests through a SOCKS proxy and logs proxying requests using the existing Apache logging mechanisms. It is most suitable for small workgroups that already have an Apache server installed.

Blue Coat Systems:

<http://www.bluecoat.com/>

Blue Coat Systems (formerly CacheFlow) offer a range of security gateway products that provide proxy caching with integrated bandwidth management and security facilities – e.g. web content filtering using SurfControl and virus scanning of web objects using the Symantec and Trend Micro systems.

Cisco Application and Content Network System (ACNS):

<http://www.cisco.com/en/US/products/sw/conntsw/ps491/index.html>

ACNS is Cisco's replacement for its end of life Cache Engine products. It includes caching as part of an integrated system of content management products for an Enterprise Content Delivery Network (ECDN). These products are aimed squarely at large enterprises and service providers so in the research and education community will probably be of more interest to Regional Network operators and large or geographically dispersed institutions.

LogiSense EngageIP:

http://www.logisense.com/cache_home.html

This cache server can run under Windows NT/2000/XP or Linux and is also available as a cache server appliance. It has a web based Graphical User Interface (GUI) for remote monitoring and management of the cache, supports plug-ins to enhance the functionality of the server (including a content filtering plug-in for the Cerberian Internet Access Control software), caches Domain Name Server (DNS) entries, supports Web Cache Control/Coordination Protocol (WCCP) version 1 and 2, has the ability to be monitored using SNMP and authenticates users against a RADIUS server or Windows NT LANMANAGER Protocol (NTLM).

Microsoft Internet Security and Acceleration (ISA) Server:

<http://www.microsoft.com/isaserver/>

ISA Server replaces the previous Windows NT based Microsoft Proxy Server (which has been discontinued and will be unsupported come January 2004). As the name suggests it is more than just a proxy web cache; it also provides firewalling and web acceleration features as well. It comes in two editions; Standard and Enterprise. Standard is aimed at workgroups and small businesses whereas enterprise is targeted at the centralised management of large networks and can support clusters of firewalling, accelerating and caching servers.

Microsoft is partnering with a number of companies to extend the ISA Server technologies, e.g. Venation are offering an integrated pre-fetching subsystem which accelerates popular sites by opportunistically caching their content:

<http://www.venation.com/>

Network Appliance NetCache:

http://www.netapp.com/products/netcache/netcache_family.html

Network Appliance have a range of hardware based cache appliances that include high availability features and support for less commonly cached protocols such as the Usenet News NNTP protocol and various streaming media formats.

Novell Volera Excelerator:

<http://www.novell.com/products/volera/>

The Novell Volera Excelerator was originally called the Novell Internet Caching System (ICS). It can be used as either a caching proxy in front of a group of clients or can provide web site acceleration. Volera has also been sold as an Original Equipment Manufacturer (OEM) product to a number of other vendors, including IBM that released it as the Netfinity server on IBM hardware.

Squid:

<http://www.squid-cache.org/>

Squid is a freely available caching web proxy, targeted at Unix/Linux hosts (though there is a version that also runs on Windows boxes). As well as proxying for normal HTTP, FTP and Gopher sessions it can proxy SSL, can take part in cache hierarchies, implements Internet Cache Protocol (ICP), Hypertext Caching Protocol (HTCP), Cache Array Routing Protocol (CARP), Cache Digests and WCCP, can provide HTTP server acceleration and caching of DNS lookups, has SNMP management facilities and supports a wide variety of access controls. Many new ideas in web caching get implemented in Squid first as it is the dominant, open source platform on the market. Squid is also found inside some commercial cache appliances.

10.2 Web Servers

Apache mod_bandwidth:

http://www.cohprog.com/mod_bandwidth.html

This Apache module provides the facility to set server-wide or per-connection bandwidth limits based on factors such as the directory being accessed, the size of the file being retrieved and/or the remote IP address or domain name.

Apache mod_throttle:

http://www.snert.com/Software/mod_throttle/

This module can also perform bandwidth limiting. It can throttle used bandwidth based on the remote IP address of the requesting client; an authenticated username from the remote client, the local user that owns the requested file, and/or the directory, location or virtual server that contains the requested file. This can be used in a variety of different ways, for example on a home page server to provide a maximum allowable bandwidth over a defined period to prevent suddenly popular web pages from soaking up all the available bandwidth - the so called 'Slashdot Effect', named after a certain website (<http://slashdot.org/>) that has a tendency to cause this to happen on a regular basis.

Arterial Software aXesW3 web server accelerator:

http://www.arterialsoftware.com/products/w3_advantages_index.html

The accelerator contains bandwidth throttling facilities that can be applied to an entire site or particular requests. Combined with its on-the-fly compression (see below) it allows individual or groups of URLs to be given a fixed maximum byte rate.

Microsoft Internet Information Server (IIS):

<http://www.microsoft.com/iis/>

and:

http://www.microsoft.com/technet/treeview/default.asp?url=/technet/prodtechnol/windowsserver2003/proddocs/standard/qos_throttlebw.asp

Microsoft IIS web server includes a bandwidth throttling control using the Windows Packet Scheduler. It has a minimum bandwidth of 8Kbit/s due to the Packet Scheduler limitations. If an IIS server is hosting more than one web site, each site can have a different throttle level applied.

Zeus Client Bandwidth Module:

http://www.zeus.com/products/zws/modules/client_bw.html

This is an optional module for the Zeus web server that allows bandwidth to be throttled based on the Multipurpose Internet Mail Extensions (MIME) type of the resource being returned, the size of the resource, an authenticated remote username, or the IP address of the requesting client. It can work in a clustered environment by sharing usage information amongst all the servers in the cluster using a shared file system such as Network File System (NFS).

10.3 Rate Limiting and Packet Shaping

Allot Communications:

<http://www.allot.com/>

The Allot NetEnforcer product line is a range of Linux based appliances offering bandwidth management features with integrated access controls, including the ability to block objects based on their content. The Allot NetAccountant software provides statistics and reporting on network usage, with extensive facilities for accounting and billing.

FortiNet:

<http://www.fortinet.com/>

FortiNet's FortiGate security appliances offer a range of ASIC assisted services integrated in a single box - including firewalling, real-time anti-virus scanning, network intrusion detection/prevention and bandwidth management.

Linux Advanced Routing:

<http://lartc.org/>

The Linux kernel comes equipped with a set of extensible and highly sophisticated bandwidth management tools, comparable to high end

commercial bandwidth management systems. It can use multiple queues with both classful and classless queue disciplines, can classify, do Layer 7 application snooping for identifying flows, and policy based routing based on packet classifications.

Packeteer:

<http://www.packeteer.com/>

Packeteer offer a range of bandwidth and application performance management hardware and software tool. A Packeteer based solution allows a network manager to monitor what traffic is actually passing over the network (using PacketSeeker) and then control the bandwidth usage of the various applications (using PacketShaper). The system performs Layer 7 application data snooping on the packets that traverse the Packeteer box and so can handle protocols such as Kazaa.

Throttled:

<http://allmacintosh.xs4all.nl/preview/278580.html>

Designed for MacOS X, this piece of software throttles traffic heading for the Internet from the Mac whilst leaving LAN traffic untouched.

10.4 Quality of Service

Cisco Switches and Routers QoS Tools:

http://www.cisco.com/univercd/cc/td/doc/cisintwk/ito_doc/qos.htm#wp1024961

Many Cisco switch and router products provide QoS tools, although these are not available in all Feature Sets. Their routers in particular offer many options for the identification and classification of packets that can then be applied to a variety of queuing mechanisms. Recent versions of Cisco's Internet Operating System (IOS) include Network-Based Application Recognition (NBAR), which provides Layer 7 snooping into packets in order to identify and mark flows more accurately.

Linux DiffServ Tools:

<http://diffserv.sourceforge.net/>

The Linux kernel packet classification and filter tools also feature the ability to mark and reclassify packets. Linux has had QoS tools such as RSVP available for many years and the operating system is used in both end points (hosts) and intermediate nodes (routers). Several large network vendors are already looking at using Linux in some of their future products so QoS support is likely to expand further in the next few years.

10.5 Compression

Apache mod_gzip:

<http://sourceforge.net/projects/mod-gzip/>

The mod_gzip module provides an Internet Content Acceleration (ICA) function for Apache. It compresses suitable content on the fly as it is delivered to the client and requires no special client software.

Arterial Software aXesW3 web server accelerator:

http://www.arterialsoftware.com/products/w3_index.html

This is a web server accelerator that makes use of idle Central Processing Units (CPU) cycles on modern web servers to compress data before sending it. It can actually reduce the overall CPU load for SSL encrypted sites as the compression algorithms can take fewer cycles to compress the data than the SSL algorithm takes to encrypt it. By compressing the data prior to encryption, less data has to pass through the SSL algorithm and so fewer overall CPU cycles are required.

Packeteer PacketShaper Xpress:

<http://www.packeteer.com/prod-sol/products/xpress.cfm>

This is a software upgrade to the Packeteer PacketShaper traffic shaper that provides compression based acceleration features.

10.6 Filtering

Cleanfeed:

<http://www.bofh.it/~md/cleanfeed/>

Cleanfeed provides spam filtering for Usenet transit servers. As well as scanning incoming newsfeeds for spam and telling the server to reject any that it finds, it can also block binary postings to non-binary groups and discard HTML postings.

DansGuardian:

<http://dansguardian.org/>

DansGuardian is a free, Unix based web content filtering system. It provides URL filtering/blocking, and can filter the content of text documents (including phrase matching). By default it is set up to content filter for young children but the level and amount of filtering is under complete control of the administrator.

Microsoft ISA server:

<http://www.microsoft.com/isaserver/>

Microsoft ISA can provide filtering for both web and e-mail traffic. As well as HTTP, FTP and SMTP protocols, it can also filter H.323 media streams and Remote Procedure Calls (RPCs). Microsoft state that its RPC filtering provides protection for Outlook users when talking to an Exchange server without the use of a Virtual Private Network (VPN).

SquidGuard:

<http://www.squidguard.org/>

SquidGuard is a combined filter, redirector and access controller for the Squid caching web proxy. It is free, extremely fast, and gives the administrator full control over blacklisted servers/URLs. Blocking can be done with powerful regular expression matching rules whilst requests for blocked resources can be redirected to an information page. Banner advertisements can be replaced with empty images. Access rules vary depending on the current time/date/user group, whilst unregistered users can be redirected to a registration form. SquidGuard does not however, filter text or code inside documents.

SurfControl:

<http://www.surfcontrol.com/>

SurfControl sell a number of filtering systems, including server based solutions for web content filtering, e-mail filtering and instant messenger filtering. SurfControl also market CyberPatrol, which is a host-based filtering system aimed at parental and school filtering of content unsuitable for minors.

10.7 Access Control

Apache mod_access:

http://httpd.apache.org/docs/mod/mod_access.html

This module provides directives in the Apache configuration file to allow access control lists to be constructed based on client hostname; client IP address, the client's request, and the method used. The mod_access module is a fundamental module in the Apache HTTP server and is shipped with it.

Cisco IOS switch/router ACLs:

http://www.ja.net/CERT/JANET-CERT/prevention/cisco/cisco_acls.html

Cisco IOS ACLs can be used to permit or deny different types of traffic based on the protocol in use, source or destination addresses and port numbers.

They are very useful as a quick means of providing simple firewalling of protocols at both border and internal routers on a site.

Microsoft IIS:

<http://www.microsoft.com/iis/>

The IIS web server can provide access controls to the content that it serves based on authenticated remote users, the remote IP address or domain name of the client browser.

11. Costs, Popularity and Relevance of Products to FE/HE

11.1 Caching

There are both commercial and free software packages available for web caching. The Squid web proxy cache is free, widely used in education and research, and was used to form the JANET Web Cache Service. It is still being actively developed and is shipped with many Linux distributions. There is a large community of experienced Squid administrators, especially in HE.

Microsoft ISA server also enjoys popularity within the UK academic community. ISA not only provides caching infrastructure but also comes with firewalling, filtering, and web acceleration features. Being a Microsoft product it integrates well with existing Microsoft based server environments.

11.2 Rate Limiting and Packet Shaping

The bandwidth throttling modules and controls available for Apache and IIS are widely deployed and used. Many sites will already have these available, even if they do not use them at the moment. As these modules are often free once the web server is installed, it makes sense to investigate how they can be used to help as part of an overall bandwidth management strategy.

Packet and traffic shaping appliances can potentially save institutions a large amount of money on bandwidth upgrades. Commercial products however are often priced based on the bandwidth that they can handle and so can quickly price themselves out of the market for HE sites where 100Mbit/s links to a nearby Regional Network and internal gigabit networks are common. They may be of more use to FE sites with lower bandwidth links. Not only will they help make better use of a comparatively scarce external bandwidth resource but they will also be more reasonably priced for typical FE institution external connection bandwidths of between 2Mbit/s and 10Mbit/s.

11.3 Quality of Service

Most of the mainstream network hardware vendors now ship products that have provision for QoS management. RSVP may be useful for providing guaranteed bandwidth to both video conferencing applications and GRID computing initiatives.

11.4 Compression

Compression on the fly is built in or freely available for many of the most commonly used web server products. For a high volume web server (such as the main web site for an institution) the bandwidth savings realised by opportunistic compression may be substantial. If nothing else it also improves

the end user experience as a compressed object obviously takes less time to download over slow modem links.

11.5 Content Filtering

Content filtering is widely deployed in the UK education and research community. Many sites have e-mail filters (both for spam and viruses) in place and an increasing number (especially in FE) have web traffic content filtering. Web filtering is often rolled into the cost of the proxy caching solution, which can make it cost effective.

11.6 Access Control Lists

Access controls are typically an integrated feature in both web servers and network devices such as switches and routers. There should therefore be no additional costs for their use aside from staff time for training; policy decisions and implementation, and system monitoring and maintenance.

12. Future Directions

Predicting the future is always difficult and is no less so when the Internet is involved. Protocols and technologies appear and achieve widespread use at an alarming rate and much of what is now commonplace was unheard of or only in research laboratories five years ago. Nevertheless what follows are some possible pointers for future bandwidth management issues and solutions.

Firstly, Peer-to-Peer (P2P) filesharing currently shows little sign of losing its popularity. Education and research institutions need to be aware of the potential of filesharing applications to unleash large amounts of potentially copyright infringing traffic onto their networks.

Content Distribution Networks are already used by large commercial corporations to ensure that end users have a copy of the much of the company's output located on a server topologically closer to them than the main company's web server. These may prove to be useful for academic sites, especially if high bandwidth media streams are likely to be downloaded - multimedia course work for distance learners, video of graduation ceremonies and important press releases or research results, for example.

Tell us what you think

Technical Guides are produced and published by UKERNA for use within the JANET Community. We welcome your comments on all aspects of this document and on any other UKERNA publication. Please direct feedback to JANET Customer Service, at the address below, or e-mail us directly at:

service@janet.ac.uk

UKERNA manages the networking programme on behalf of the higher and further education and research community in the United Kingdom. JANET, the United Kingdom's education and research network, is funded by the Joint Information Systems Committee (JISC).

For further information please contact:

JANET Customer Service
UKERNA
Atlas Centre, Chilton, Didcot
Oxfordshire, OX11 0QS

Tel: +44 (0) 1235 822 212
Fax: +44 (0) 1235 822 397
E-mail: service@janet.ac.uk

Copyright:

This document is copyright The JNT Association trading as UKERNA. Parts of it, as appropriate, may be freely copied and incorporated unaltered into another document unless produced for commercial gain, subject to the source being appropriately acknowledged and the copyright preserved. The reproduction of logos without permission is expressly forbidden. Permission should be sought from JANET Customer Service.

Trademarks:

JANET®, SuperJANET® and UKERNA® are registered trademarks of the Higher Education Funding Councils for England, Scotland and Wales. The JNT Association is the registered user of these trademarks. All other trademarks are property of their respective owners.

Disclaimer:

The information contained herein is believed to be correct at the time of issue, but no liability can be accepted for any inaccuracies. The reader is reminded that changes may have taken place since issue, particularly in rapidly changing areas such as internet addressing, and consequently URLs and e-mail addresses should be used with caution. The JNT Association cannot accept any responsibility for any loss or damage resulting from the use of the material contained herein.

Availability:

Further copies of this document may be obtained from JANET Customer Service at the above address. This document is also available electronically from:

<http://www.ja.net/documents/>
