MINT 709 – Internetworking Project

Bandwidth Management

Submitted : August 14, 2006

By

Jason W. Fernyc

Professor M.H. MacGregor
First Reader

Professor J.N. Amaral
Second Reader

ABSTRACT

Traffic prioritization and Quality of Service (QoS) are important components in the offering of network services and have enhanced importance when converged traffic transits the network. Each traffic type may place unique requirements on the network infrastructure.

While this functionality can provide benefits to the network operator, barriers to implementation such as cost and complexity often preclude the Small Internet Service Provider (ISP) from pursuing this important functionality with their networks.

This paper provides a compilation and analysis of the methods available to small/medium networks to provide a viable Quality of Service (QoS) implementation to optimize traffic flow in these architectures.

This paper explores bandwidth management techniques available in the marketplace, and compares them against the requirements of the target audience. A recommendation for a bandwidth management solution is developed based upon meeting these requirements, including: low Capital Expenditure (CapEx), low Operations Expenditure (OpEx), along with low complexity to implement.

Equally important to the core requirements above, is that the solution proposed can be deployed independent of the network equipment vendor that an ISP is utilizing. It can also be deployed independent of a specific Physical layer (PHY layer) technology. The solution proposed provides maximum ability to migrate to future developments/technologies that vendors may put forth.

Table of Contents

1.0 Introduction

Bandwidth management with a network is an important process to optimize performance and can be utilized to manage some of the costs associated with operating an organization's network. The use of bandwidth management is explored within this paper, as a proposed solution applicable to small/medium sized networks to optimize their network performance, reduce costs, and offer superior quality of service to their subscriber base with minimal investment.

Small/Medium sized networks may be found in areas such as: industrial networks set-up on a temporary basis, small educational institutions, and also the main focus audience of this paper: small to mid-sized Internet Service Providers (ISPs).

This last grouping of network type, Small ISP, can benefit significantly in both monetary and efficiency terms by adding some bandwidth management functionality to their existing systems. This functionality is lacking in most of the networks within this sector due to the potential costs and complexity to implement the mainstream solutions available in the marketplace.

This paper explores several options available to implement bandwidth management functionality, and proposes a method that is both minimal cost and straightforward for the target audience to implement.

The convergence of traffic types on modern networks spans even those networks of the small ISP provider. Traffic flows may contain data, voice or video, which drives a requirement to develop bandwidth management schemes within networks of any size in order to maximize the success of the co-existence of these traffic types.

The tolerance of each traffic type outlined above to such parameters as: delay, jitter and packet loss differs. Each of these parameters translates to directly impact the user's perceived experience.

2.0 Traffic Profiling

Today's converged networks carry widely varied types of traffic, each with their own characteristics which place unique demands upon the network. To accommodate these distinct characteristics and requirements it is useful to first implement a classification methodology that allows for the identification of traffic into groups of similar characteristics, or to identify any traffic deemed more important than generic traffic.

Once these traffic groups and types can be identified, then a bandwidth management implementation can provide separate handling for these traffic types. This handling may be a prioritization, or a de-prioritization of the specific traffic.

A custom set of classification rules can be developed for each group and type of traffic in order to support protocols and applications according to local policies. These rules can also be modified and adapted as desired moving forward.

2.1 Optimization Parameters

The main objectives of implementing bandwidth management are varied with a dependency upon the network in which such a mechanism is deemed warranted. Prioritization is a key component of a bandwidth management solution. It can be implemented as a means to optimize network performance in the following important areas across a wide range of networks:

I. To minimize the costs of providing bandwidth
II. To maximize system throughput
III. To minimize latency through the network
IV. To ensure Quality of Service (QoS) to high priority end-points

2.1.1 Cost Optimization

The demarcation point from the local area network (LAN) to the wide area network (WAN) is the point within the system where costs to carry network traffic may become significant. Often, there are different entities that control the WAN connections such as: different groups within a single organization or they may be several separate companies.

In the case of the target audience discussed within this paper, these WAN providing entities are third party transport providers such as Telus[1] or Sprint[2].

The costs involved for the transmission of data may be architected on both a bits-per-unit time (bps throughput) basis, and also on a bits-per-billing-period (total bandwidth).

The solution proposed within this paper could be applied to mitigate the amount of traffic that reaches the WAN link, and thus minimize the transport costs involved. This cost savings could be realized in one of two ways:

I. by managing which traffic streams reach the WAN connection, there is a real-time savings if the provider is charging on a per-bit basis, or
II. By managing how traffic reaches the WAN connection via prioritization, the ISP can avoid upgrading their bandwidth contract to a higher speed unnecessarily or prematurely.

A special case is that of undesirable traffic, which may consist of traffic generated via computer virus, denial of service (DoS) or non-business critical (P2P) means. This traffic type may be throttled as required to further optimize what traffic reaches a metered WAN connection, or to conserve throughput for higher priority flows to transit the WAN connection.

2.1.2 System Throughput Optimization

Throughput of traffic through the network may be optimized if signaling traffic types receive priority. Generally, these traffic types are associated with applications that depend upon connection-oriented operation. This includes any application protocols based upon the layer 4 Transmission Control Protocol (TCP).

Applications utilizing this common protocol share an important commonality, which is the dependency upon two parameters necessary for transmission:

I. SYNchronization (SYN) transmissions in the forward path to receiver
II. ACKnowledgement (ACK) transmissions in the return path to the sender

By optimizing the ability for these two signaling traffic types to transit a given network, via priority marking and processing, total network throughput can be increased over the default state of the network[3].

By allowing each of these traffic types to transit the provider's network with priority and appear at the WAN connection ahead of other traffic, the aggregate throughput achieved by users of these TCP based protocols will be higher than in a default network configuration.

2.1.3 Latency/Jitter/Packet Loss Optimization

While traditional data-only network traffic is inherently very tolerant of delay, more advanced traffic flows require real-time delivery of their information streams with a sensitivity to delay, variation of delay (jitter) and (packet) loss.

A common-place example of a traffic flow meeting this criteria is packetized voice over internet protocol (VoIP). While the exact characteristics of VoIP traffic flows depend upon the configuration utilized within a given network, there are the following common characteristics evident of this type of traffic:

I. Small packet size: typical rate of sampling sufficient to represent voice runs up to 64kbps. Some coders:decoders (codecs) perform at levels down to 8kbps. In either case, the voice samples yield a very small payload size.

II. Frequent packets generation: typical packetization of a 64kbps voice stream results in a complete packet every 20 milliseconds. This yields a net packets-per-second (pps) stream of 50pps in each direction, per phone call.

The product of small packet size, and rate of packet generation yields an effective traffic rate of approximately 100kbps (including overhead to the IP layer) in each direction for a voice over IP call.

These calculations assume the use of the ubiquitous G.711[4] codec. This codec has been selected as the coding of choice due to its ability to mimic the performance and quality of the incumbent Public switched Telephone Network (PSTN). Canadian broadband access companies such as Shaw[5], Rogers[6], and Cogeco[7] are all utilizing this implementation in their digital voice consumer offerings.

While a flow, with these characteristics, represents a relatively low volume in today's multi-megabyte-per-second (Mbps) networks, these same flows have an extremely limited tolerance for delay, jitter and packet loss. It thus becomes important to ensure that a method of handling these flows with minimal latency, jitter and packet loss is implemented. Such methods are proposed in subsequent sections of this paper.

With these characteristics optimized, the result will be a high quality calling experience, where the use of VoIP technology mirrors the Incumbent local exchange carriers[8] (ILECs) to the end users.

2.1.4 Quality of Service (QoS) - High Priority End-points

Within networks there are stations that require, or in most cases more-so desire, varying prioritized access to the general network and the resources located across the network infrastructure.

Within an organization, the rules for this prioritized access, may simply follow the structure of the organization, with priority assigned on a per-department basis, or a per-function basis. Executives may demand higher priority than functional departments lower in the organization structure.

Prioritization rule-sets may be developed in order to service higher priority stations/groups of stations, in order to meet their quality of service requirements. The end-goal being to provide these specific stations on the network with a service level of predictable and acceptable access, even during periods of network congestion.

The bandwidth management solution proposed in this paper can provide for quality of service to specific end-points, however for the target audience a more optimized method of traffic management is proposed based upon flow identification and handling.

2.1.5 Quality of Service (QoS) -High Priority Applications

Within networks there are applications that require varying responsiveness from the network in order to perform their core function(s) most effectively. Examples of traffic flows that expect a high quality of service are: voice over IP and Secure Shell (SSH).

Each of these applications require real-time or near real-time service in order to translate into a user experience that is deemed acceptable vs. expectations. Their tolerance for delay, jitter, and loss are minimal at all times.

Applications such as FTP or HTTP browsing can co-exist well on the network with lower quality of service treatment, and without explicit handling. Their tolerance for these same network impairments is much more substantial.

With a bandwidth management solution in place, real-time and near real-time applications will be able to transit the network with precedence over the other traffic types that are competing for network resources, and thus maintain a high quality of service.

If these applications were instead to transit the network via normal, or best-effort, mechanisms they will become impaired as network loading fluctuates. When the network is severely impaired, or congested, the user experience will degrade to an unacceptable level.

If the congestion, and thus the degradation, is frequent enough, the user population will move to abandon the use of these applications as they will deem them unreliable. This has implications for the ISP from a customer satisfaction perspective. If severe enough, the user-base may seek an alternate provider.

3.0 Bandwidth Management Techniques

Quality of Service (QoS) within a network depends on traffic policing, queuing and scheduling algorithms.

The policing function can be implemented independently of the access network technology in use. This allows for a common policing function to be leveraged across diverse network implementations, and achieve a cost effective means of providing a high quality of service within the network, rather than providing via significant expansion of network infrastructure, and the resulting increases in capital spending.

This paper discusses an implementation of traffic policing/control, with the design goal to be agnostic to the physical (PHY) layer implementation of the network, as well as, support multiple vendor environments. This allows for the implementation to be deployed upon commonly found infrastructure such as: xDSL[9], coaxial broadband[10], fixed wireless[11], as well as, within Satellite Communication (SATCOM)[12] environments. It also allows for deployment without dependencies upon associated vendor features of core networking companies such as: Cisco[13], HP[14], 3COM[15] and Nortel[16].

By proposing a solution that is independent of the PHY layer network technology, operators are provided with flexibility for future migration and provided the ability to mix deployed technologies without material impact to the bandwidth management strategy already established.

Ultimately, the goals with this approach are to minimize any downtime to subscribers when the network topology or equipment evolves to meet future demands.

3.1 Key Network Metrics

The bandwidth management techniques, researched and presented within this paper are designed to offer improvement in important network traffic metrics such as latency, jitter and packet loss.

The improvement in these metrics will be measurable in the implementation phase of this project, as well as, the ultimate placement upon production networks of the target audience.

The main parameters considered within this paper are:

I. Latency
II. Jitter
III. Packet Loss
IV. System Throughput

3.1.1 Latency

Latency is the measure of delay within a transmission path[17]. This is typically measured as end-to-end delay, or round-trip-time(RTT) in terms of tens or hundreds of milliseconds for a packet to reach the destination from it's source. This parameter is generally not a factor in applications utilizing bulk data transfer, instead it becomes a significant factor when considering real-time, or delay sensitive traffic flows.

Ideally, these flows should be prioritized for delivery above other forms of traffic, to ensure that any latency on these packets is minimized. This will in turn ensure that the resulting user experience based upon these real-time traffic streams is a positive one.

In a network with default configuration, meaning no bandwidth management mechanism, real-time traffic flows will experience an increase in latency as network load grows. These flows will simply be interleaved between other concurrent traffic types.

At the point of network congestion, latency may reach levels that render the real-time streams unusable to the end-user and/or end-user application. With the addition of a bandwidth management implementation, real-time traffic will be able to be detected, marked and transit the network with priority over other flows, thus reducing the effects of congestion on latency.

The acceptable latency metric required of today's networks, in terms of one way delay is 150 milliseconds for voice over IP networks. This is considered the maximum allowable[18] delay, before quality of the call is considered degraded from toll grade.

It should be noted that the prioritization implementation proposed within this paper will inherently add a small delay to the traffic streams, beyond serialization and propagation delays. This delay is comprised of the time period to evaluate and process each packet by the policing device/application before it is passed onto the network. The delay induced is expected to be minimal when compared to the delay characteristics that these flows would be subjected to in a congested network without prioritization.

3.1.2 Jitter

Jitter is the measure of the variation of the delay parameter discussed above within a transmission path[19]. Jitter is commonly measured as an end-to-end metric, once again in units of milliseconds. This parameter can be destructive to the quality of voice streams, as the resulting voice playback will be disjoint, and choppy to the human ear.

The design goal is to keep this parameter to a minimum, which is especially important for real-time traffic. Some end-points have integrated functionality in the form of specialized de-jitter buffers to deal with this network impairment, however prioritization can be utilized to significantly reduce jitter within the network for the appropriate traffic streams.

The acceptable jitter metrics required of today's networks, specify a maximum 5 milliseconds as tolerable by most telecommunications providers[20].

In a network with default configuration, it would be expected that real-time traffic flows would experience an increase in latency, along with a corresponding increased variation in delay(jitter) as network load grows towards the point of congestion.

With the addition of a prioritization implementation for the real-time traffic, it would ensure that this traffic instead transits the network in a more predictable manner, "in-front" of other packets. This would minimize any jitter experienced as these packets transit the network.

3.1.3 Packet loss

Packet loss is the measure of packets not arriving, or otherwise being unrecoverable at the destination[21], typically represented as a percentage of total traffic. Packet loss can increase as the network load grows towards congestion. Acceptable packet loss metrics are of the order of < 1% in production networks[22] for real-time applications such as VoIP. Loss in excess of this value will result in degradation of the application/service being supported via the traffic streams.

It is possible that a packet loss may simply be perceived at the receiving station, due to the packets being queued for a significant period of time at the sending station. In this scenario the packet loss is not necessarily a function of loss over the network, but rather loss due to timeout being exceeded. Utilizing a prioritization scheme will provide improvements in these packet loss metrics as resent packets can be minimized.

3.1.4 System Throughput

Total system throughput is dependent upon many factors, including the timely handling of signaling traffic within the network which controls the bearer traffic flows. In networks, such signaling traffic typically initiates or acknowledges the continuation of a network traffic stream. By prioritizing the transmission of this traffic, the associated data streams should reach maximum throughput over the network, at a rapid rate.

A practical example of this would be the prioritization of TCP ACK traffic, which is designed to enhance throughput of TCP flows, as discussed later within this paper.

3.2 Queuing Techniques

A mechanism of sorting traffic into priority levels must be implemented at each network device. This is especially important as the network loading trends towards congestion. Packets can be identified/classified and placed in different types of queues before delivery onto the network. The queues can be variable size, and prioritized as to when they can empty themselves.

There are several types of queue implementations found in production networks. Some of the more prevalent strategies are listed and discussed below:

I. First-in, First-out (FIFO)
II. Weighted Fair Queuing (WFQ)
III. Class Based Weighted Fair Queuing (CBWFQ)
IV. Priority Queuing (PQ)
V. Low Latency Queuing (LLQ)

3.2.1 First-in, First-out (FIFO)

Using FIFO[23] queuing, the first packet to enter an interface is the first to leave that interface. Thus first packet into the system ends up being the first out., regardless of it's payload This is often the default mechanism for receiving and sending packets on a standard network interface.

Under congestion, quality of service will not be optimized for this type of implementation, as there is no ability to prioritize real-time traffic ahead of other traffic types. This is a very simplistic approach to queuing.

3.2.2 Weighted Fair Queuing (WFQ)

Weighted Fair Queuing[24] (WFQ) allows several network traffic types to coexist. It is designed to allocate bandwidth fairly amount the flows present on the system. If there are N flows within the system, then each will derive 1/Nth of the available channel.

3.2.3 Class Based Weighted Fair Queuing (CBWFQ)

Class based Weighted Fair Queuing[25] (CBWFQ) builds upon the functionality offered by WFQ, with the addition of using classes to give proportional bandwidth to the users.

This can be accomplished through use of creating classes based upon the IP Precedence settings for example. These markings can be utilized as weightings and the number of queues will depend partly upon the number of flows present in the system. Traffic can be classified based upon source address and port, destination address and port, along with IP precedence settings.

The WFQ mechanism interleaves smaller packets between the larger packets. The forwarding decision in WFQ is based on the first packet to finish entering the queue. This scheme thus favours shorter packets, as they will finish entering the queue ahead of larger packet sizes. The net result is a queue where low bandwidth traffic takes priority over high bandwidth traffic. This is suitable for applications such as VoIP, which inherently have a low bandwidth signature.

Once the priority queues have been serviced, any remaining IP traffic is weighted and queued accordingly. The weighting factor is dependent on the IP Precedence and can be calculated as in the following example:

I. Let IP precedence settings be P_1, P_2 and P_3.
II. Sum-up all (precedence settings+1) as:
   1. $(P_1 + 1)$, $(P_2 + 1)$ and $(P_3 + 1)$.

III. Each flow will get a proportion of the link bandwidth according to the following breakdown:

- P_1 will receive: $(P_1 + 1) / ((P_1 + 1) + (P_2 + 1) + (P_3 + 1))$
- P_2 will receive: $(P_2 + 1) / ((P_1 + 1) + (P_2 + 1) + (P_3 + 1))$
- P_3 will receive: $(P_3 + 1) / ((P_1 + 1) + (P_2 + 1) + (P_3 + 1))$

An inherent limitation in the design of a WFQ implementation is that higher priority traffic does not necessarily transit with true priority in one burst. Instead, as the number of flows increases, all flows receive an increasing fraction of the total channel.

Higher priority flows still transit with precedence however the impact of this precedence is diminished in comparison to when there were a smaller number of flows. The reason for this is that as the number of flows increases, it will drive the denominator up significantly. This in turn will lower the sensitivity of the system to any one flow.

3.2.4 Priority Queuing (PQ)

A priority queuing[26] (PQ) implementation typically utilizes four queues: High, Medium, Normal and Low. Each packet will be directed into the appropriate queue from this list, depending upon its classification. Any packets that cannot be classified are placed into the 'normal' queue. This normal queue is sometimes referred to as the best effort queue.

In PQ operation, higher priority queues must be emptied before the lower priority queues can obtain service. A drawback of this characteristic is that lower priority queues have a potential to be serviced at a much lower rate than they need to empty. This phenomenon is called Starvation.

Depending upon the volume of high priority traffic in the system, lower priority queues could receive no service at all for an extended period of time. In this case starvation within the system would be driven to high levels and this would impair the traffic residing in the lower priority queues by a significant factor.

3.2.5 Low Latency Queuing (LLQ)

Low latency Queuing[27] (LLQ) consists of a single priority queue implementation, which brings strict priority queuing to Weighted Fair Queuing (WFQ). The strict priority queuing component allows for delay-sensitive information (aka real-time traffic), to be processed and transmitted first.

This effectively provides the real-time traffic flows with preferential treatment over other traffic flows in the system. This deals with TCP exchange very well and is especially suitable to real-time traffic, such as Voice over IP(VoIP) flows. Jitter is minimized in a system utilizing a LLQ implementation.

4.0 Prioritization within Practical Networks

In order to implement traffic prioritization within a network, it must first be determined which parameters will be utilized for classification of traffic. Each type of network traffic inherently has a distinct set of characteristics, or signature, that can be leveraged for this classification function. Common parameters than can be utilized as classifiers are:

I. source IP address
II. destination IP address
III. source port
IV. destination port
V. header flag settings

4.1 Detectable Classifiers

Each packet that transits the network has multiple fields that can be utilized to classify the traffic as belonging to a specific traffic flow. These fields may be addresses of stations, addresses of applications (ports) or markings within the header information designating a specific packet as requiring special servicing.

The bandwidth management solutions proposed within this paper can optimize the allocation of bandwidth within the system by implementing all of these function parameters.

4.1.1 Source IP Address

Traffic classifiers may be developed that are based upon the source address, or group of source addresses (e.g. subnet) where network traffic originates. The purpose of this method would be to ensure that traffic from a specific station, or group of stations, could be detected, and classified as a separate traffic type from all other traffic within the network. Once classified it could then be handled according to specific rule-sets defined alongside the classification rules. The end result would be that these stations could then receive prioritized access to the network per the configured rule set and hierarchy. In this model, the priority provided would be granted, regardless of the type of traffic these stations are sourcing.

Configuration of this classifier would look for a match to the layer 3 IP address configured on the target station.

4.1.2 Destination IP Address

Traffic classifiers may be developed based upon the destination address, or group of destination addresses (e.g. subnet) where traffic is destined. As with the source address method defined above, the purpose of this method would be to ensure that a specific station, or group of stations, should always maintain priority over the peer population, in this case with an emphasis on inbound traffic. Configuration of this classifier would typically look for a match to the layer 3 IP address configured on the target station.

4.1.3 Source Port

Traffic classifiers may be developed based upon the layer 4 source port, or group of source ports (e.g. port range) where traffic originates. The purpose of this method is to ensure that a specific application, or group of applications on the local network, will maintain priority over the peer population, regardless of which stations they are originate from. Configuration of this classifier would typically look for a match to the layer 4 Transport port number utilized by the application(s).

When classifying based upon ports, the destination port classifier strategy, as detailed in the next section, is a more applicable method as many applications pick an arbitrary source port rather than a fixed port.

4.1.4 Destination Port

Traffic classifiers may be developed based upon the layer 4 destination port, or group of destination ports (e.g. port range) where traffic is received. The purpose of this method is to ensure that traffic inbound to a specific application, or group of applications, will maintain priority over the peer population, regardless of which stations they are received by. Configuration of this classifier would typically look for a match to the layer 4 Transport port number utilized by the application(s).

4.1.5 Header Flags

Traffic classifiers may also be developed based upon the settings of the header flags within a given packet. These flags may exist at layer 3 - IP header, or layer 4 - TCP header. The purpose of this method is to ensure that packets bearing a specific signature within their headers should always maintain priority over the peer population, regardless of which stations they are sent or received by, and regardless of which application/traffic is being transported within the packets.

In exploring each of the classification techniques discussed above, three general categories for implementation are considered:

I. Flow-Based Model
II. Station-Based Model
III. Hybrid-Based Model

4.2 Flow-Based Model

Consists of a rule-set within the bandwidth management device that will prioritize traffic on a per-flow basis regardless of the stations involved in sourcing, or receiving the various traffic streams.

Traffic control utilizing this model ensures that each application within the network can receive a dedicated level of quality of service and handling. This model handles all types of traffic, and handles real-time traffic very well.

Applications with equal priority parameters will obtain an equal share of network resources. In the presence of congestion, higher priority flows will be able to transit the network, whereas lower priority traffic flows will follow only after high priority flows have been serviced.

4.2.1 Application Classification

Flows within a network can be evaluated on the basis of the layer 4 port numbers[28] in order to determine which applications they are associated with.

These signatures allow classification of traffic on the basis of application. These applications can then be grouped according to the quality of service requirements they present to the network. This grouping process is illustrated in the sections below:

4.2.1.1 Real-Time Application Traffic

Voice over IP (VoIP) is the most prevalent real-time traffic type in the commercial networking world at the present time. Telephony calls utilizing this traffic type require a maximum of 110kbps (using G.711[29]) codec and typical IPV4 packetization along with Ethernet framing. This bandwidth requirement is in each direction as the voice channel will generally be bi-directional and simultaneously active during steady-state operation.

There are schemes to minimize bandwidth requirements further using voice activity detection (VAD) along with lower rate vocoders (e.g. G.729 @ 8kbps), however those techniques do not see wide scale usage in the networks of the target audience to date.

This will change in the future as VoIP network mature to the point of wide-scale adoption across enterprise and residential use. VoIP traffic is primarily based upon the UDP transmission mechanism and thus the flows are connection-less in nature.

4.2.1.2 Near Real-Time Application Traffic

Second in priority to real-time traffic flows, are traffic flows that require near real-time service. Two prevalent examples of traffic that meets this classification would be telnet and Secure Shell[30] (SSH). As with the VoIP flow above, a single telnet or Secure Shell (SSH) session requires a relatively low overall bandwidth to/from the network, once again in both directions, however the protocols are very sensitive to latency and jitter.

The telnet/SSH protocols themselves are bidirectional, however the method in which they are employed(input:wait for response:input:wait for response), inherently means that only one direction will be operating at full speed, while the other direction has relatively low bandwidth requirements.

Enhancing the ability for these near real-time traffic types to transit the network over generic traffic will allow acceptable performance for the end-users of these applications, be they human or automated tools.

4.2.1.3 Best Effort Traffic

This traffic type may simply consist as a catch-all for all traffic within the network that does not successfully match one of the configured traffic classification rule sets. This traffic will be allowed through the network when all classified traffic is processed and forwarded. In times of network congestion, the service available to this traffic type will be minimal.

4.2.1.4 Nuisance Traffic

This traffic type represents network flows that are undesirable and should be prioritized below all other traffic. Examples of traffic that may match these would be denial of service (DoS)[31], peer-to-peer(P2P), or virus related traffic types.

Typically, classification of these types of traffic can be accomplished via evaluating the layer 3 and layer 4 characteristics of the traffic against a known signature profile. By creating a classifier that triggers on these signatures, and subsequently lowering the priority of this traffic below best-effort service, it will reduce the negative effects on the network that such traffic would impose in a native network configuration.

4.2.1.5 Application Classification Implementation

The figure below illustrate the classification process in taking in the default traffic stream and reordering the packets according to 3 example queues, real-time, near real-time and best effort.

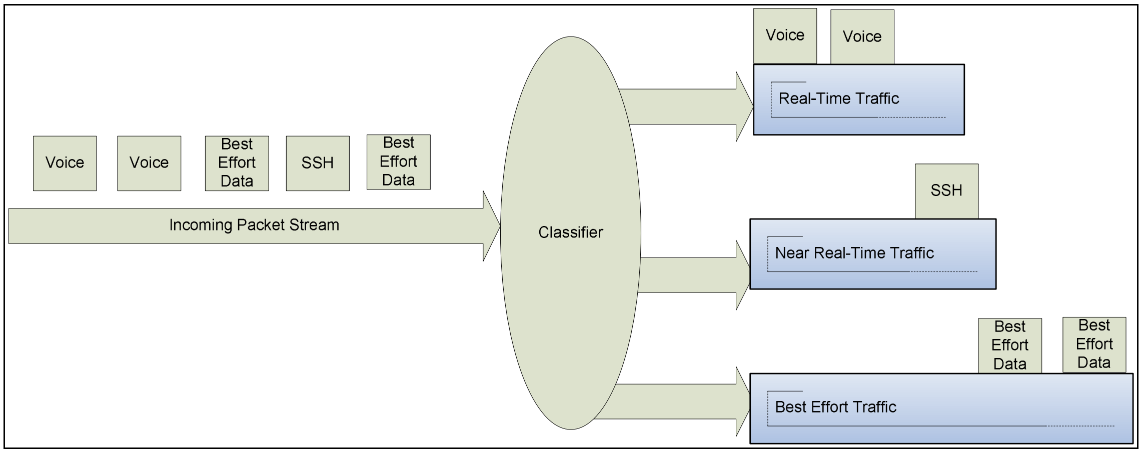

Figure I – Application Classification

4.2.2 Non-Application Classification

Flows within a network can also be evaluated on the basis of the header information transmitted in each packet. This important header information may exist at layer 3 and/or layer 4. The evaluation of each of these headers can be very useful in grouping flows for prioritization within the network.

4.2.2.1 Layer 3 Classification

The layer 3 (IP) header contains a field called Type of Service (TOS) or more recently renamed Differentiated Services Code Point (DSCP). These fields exist for the purpose of making the network aware of the relative priority that the payload contained within the layer 3 packet should receive. Either can be utilized as the basis for a classifier for traffic matching a known profile.

4.2.2.1.1 Type of Service (ToS) Field

The three bit ToS field, and it's typical usage is discussed below, beginning with a diagram illustrating their placement within a typical Internet Protocol Version 4 (IPV4) header.

This field can also be referred to as the precedence field, given that it's function is to designate whether a packet requires special handling within the network. A full discussion is not undertaken in this paper, and is instead available within RFC 1349[32].

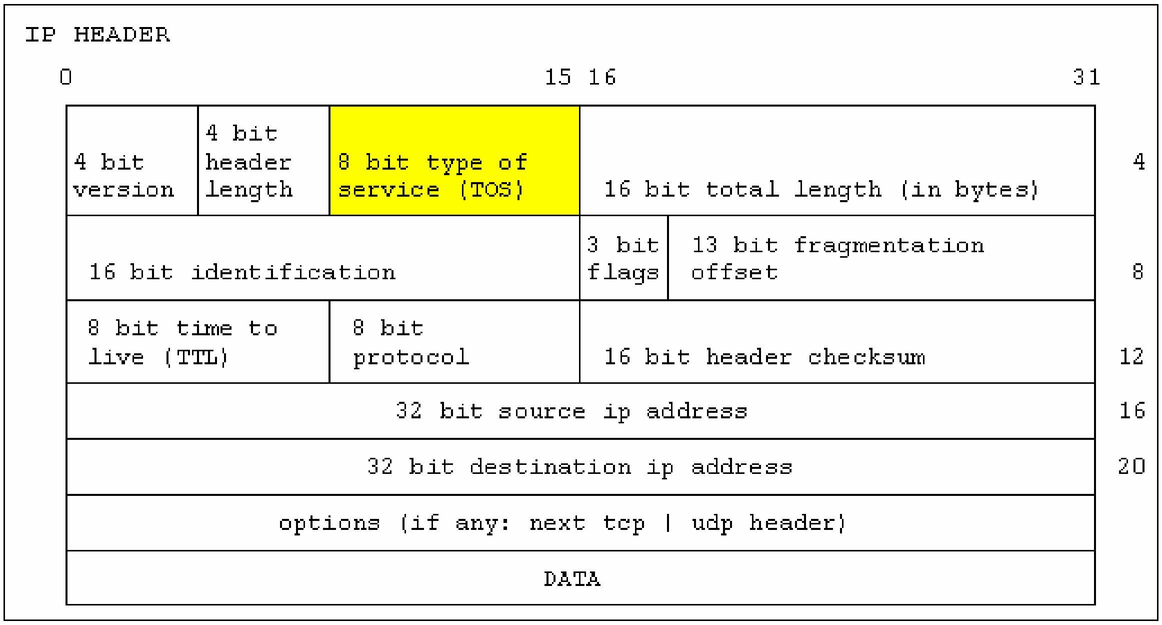

IP HEADER

0 15 16 31

4 bit version	4 bit header length	8 bit type of service (TOS)	16 bit total length (in bytes)	4
16 bit identification		3 bit flags	13 bit fragmentation offset	8
8 bit time to live (TTL)		8 bit protocol	16 bit header checksum	12
32 bit source ip address				16
32 bit destination ip address				20
options (if any: next tcp | udp header)				
DATA				

Figure II – IPV4 header, illustrating the ToS fields in Yellow

Further analysis of the field shown above in yellow, yields the following breakdown of the bits shown in figure III:

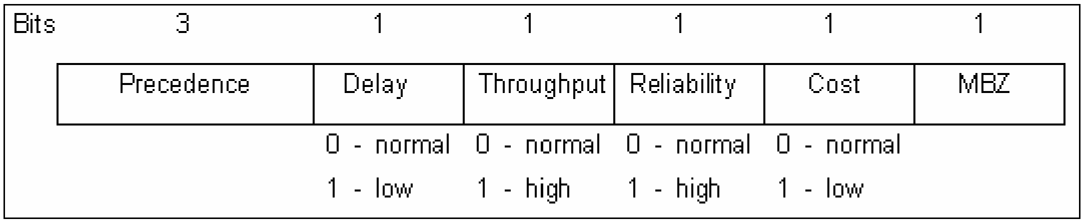

Bits	3	1	1	1	1	1
	Precedence	Delay	Throughput	Reliability	Cost	MBZ
		0 - normal	0 - normal	0 - normal	0 - normal	
		1 - low	1 - high	1 - high	1 - low	

Figure III – IPV4 header, illustrating the TOS bit positions/functions[33]

It can be seen from the diagram above that setting the TOS bits to a value of '1', is a method to mark the packet as requiring enhanced quality of service over the default case.

Defining these fields to describe their functionality:

Parameter	*Setting*	*Function*
Delay	1	Request for low delay
Throughput	1	Request for high throughput
Reliability	1	Request for high reliability
Cost	1	Low Cost

4.2.2.1.2 Differentiated Service Code Point (DSCP) Field

The Differentiated Service Code Point (DSCP) fields are a 6 bit selector that builds upon the ToS fields discussed above. RFC2474[34] fully defines the bit values within these fields, and their relationship to the Type of Service bits.

The DSCP implementation overlaps the legacy ToS (aka. Precedence) field for backward compatibility. The newer DSCP field is also twice the size of the legacy ToS field, being 6 bits long vs. the original 3 bits that ToS provided. This means that if the values of DSCP are carefully chosen then backward compatibility can be easily achieved. It also means that the granularity possible with DSCP, and thus the classifiers that could be derived, are increased significantly over the original ToS field implementation.

This results in the concept of DSCP "classes", with each class being a group of DSCPs mapping to the same Precedence value. Values within a DCSP class offer similar network services but with slight differences. Examples of this would be differing levels of service such as "gold", "silver" and "bronze" service from the same service provider.

The standardized class mapping is defined below. It can be seen that the DSCP value can be found by multiplying the precedence value by 8 :

DSCP	***Precedence***	***Purpose***
0	0	Best effort
8	1	Class 1
16	2	Class 2
24	3	Class 3

DSCP	***Precedence***	***Purpose***
32	4	Class 4
40	5	Express forwarding
48	6	Control
56	7	Control

4.2.2.2 Practical use of TOS/DSCP

Voice over IP (VoIP) providers around North America are beginning to focus more on the ToS bit capabilities they deploy this time sensitive application across their networks. This research has shown that voice signaling on many of these networks is currently being given a (TOS) priority of 3, while voice payload(bearer traffic) is given a priority of 5. These values map to DSCP values of 24 and 40 respectively. Vendors, such as Motorola, Ambit, and Cisco, who produce equipment placed at the customer premises (CPE) are hard-coding these values as defaults into their networking equipment at the request of the (VoIP) service providers.

This priority marking scheme allows for Quality of Service to be enhanced both in the 'last mile' that these providers service, and also has implications further north into the network as the traffic transits the provider's, and perhaps third party networks, on its path to the Public switched telephone network (PSTN).

Type of Service (TOS) has not been widely used despite being part of the TCP/IP protocol standard for a lengthy period of time. The feature is supported by routing protocols such as OSPF and IS-IS, however application support has lagged behind these protocols, making usage of the TOS parameter within these functions relatively limited.

The re-emergence of the IS-IS protocol, as networks move towards IPV4/IPV6 hybrid configurations, may bring about renewed focus on the TOS/DSCP fields within the routing space, while the increase in converged networks will bring about a renewed focus on these fields for more general purposes.

4.2.2.2.3 Layer 3 Classification Implementation

The figure below illustrates the classification process in taking in the default traffic stream and reordering the packets according to 3 example queues, TOS of 3, 5 and the default best effort queue.

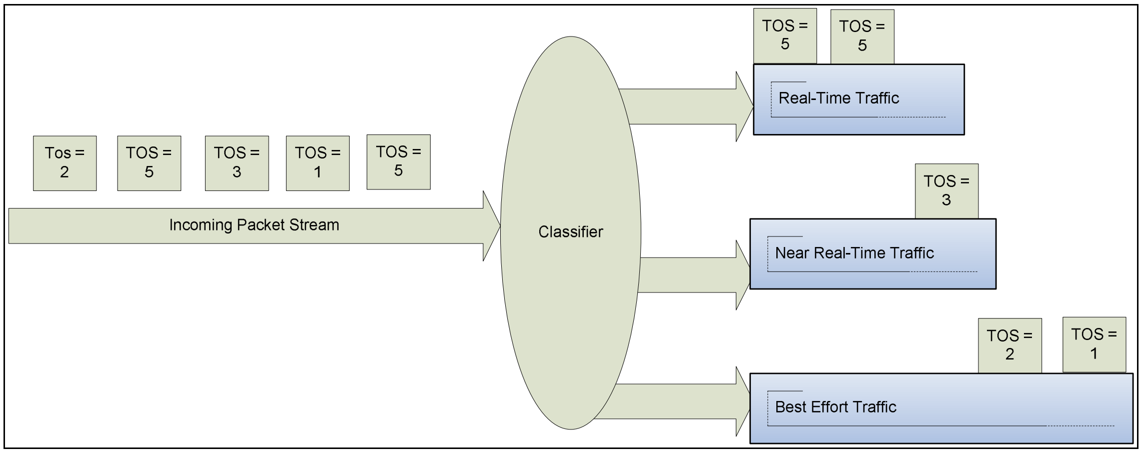

Figure IV – Layer 3 Classification

4.2.2.3 Layer 4 Classification

The layer 4 header, when utilizing the TCP connection-oriented protocol, contains several FLAGS that can be evaluated and leveraged to enhance network traffic performance.

The setting of the flags can designate whether signaling traffic for the TCP connection or payload traffic is contained within the packet. This allows the system to process each of these traffic types independently. It is proposed within this paper that this signaling traffic should receive priority.

4.2.2.3.1 TCP ACKnowledgement

An important facet of the bandwidth management and prioritization schemes recommended within the paper is the prioritization of Transmission Control Protocol ACKnowledgement (TCP ACK) packets[35].

When a network experiences saturation or congestion, downstream throughput is reduced due to delay induced in the upstream direction, which is the path that the ACK packets transit. By prioritizing the TCP ACK traffic in one direction, the throughput of the corresponding data stream can be enhanced[36].

This type of traffic has a distinct signature, which is discussed below. Such a signature is useful in the classification and subsequent prioritization techniques.

There are potential downfalls in implementing such a prioritization method on TCP ACK packets. There is a possibility that other protocols will attempt to leverage the priority placed upon the ACK packets in order to accelerate their own performance. Application developers may try to masquerade their traffic within packets with ACK flags set.

Careful analysis of the packets being classified as TCP ACKs can avoid this potential pitfall. Normal TCP ACKs will typically contain a total packet size in the order of 80 bytes or less. Any rogue protocol attempting to mimic the ACK packet in order to transport its information will typically utilize a much larger size, due to the payload contained within these non-desirable packets contain in contrast to a legitimate ACK packet. By evaluating the packet size, it is a straightforward process to recognize a non-legitimate ACK transmission.

The TCP protocol is defined within RFC793[37] Figure V illustrates the TCP header defined within this standard, with emphasis on the ACK flag within the header structure. The ACK flag is shown below in RED.

```
.-------------------------------+-------------------------------.
|          Source Port          |       Destination Port        |
|-------------------------------+-------------------------------|
|                        Sequence Number                        |
|---------------------------------------------------------------|
|                    Acknowledgment Number                      |
|-------------------+-+-+-+-+-+-+-------------------------------|
|  Data |           |U|A|P|R|S|F|                               |
| Offset| Reserved  |R|C|S|S|Y|I|            Window             |
|       |           |G|K|H|T|N|N|                               |
|-------+-----------+-+-+-+-+-+-+-------------------------------|
|           Checksum            |         Urgent Pointer        |
`---------------------------------------------------------------'
```

Figure V – TCP header, including the ACK flag in RED

The TCP ACK Flag has two possible values as follows:

Setting	***Meaning***	***Function***
0	Not Set	Packet is not a TCP ACK
1	Set	Packet is a TCP ACK

Even when a TCP connection is used to send data only in one direction (such as downloading a file through ftp), TCP acknowledgements (ACKs) must be sent in the opposite direction, or the far-end station will assume that its packets were lost and eventually retransmit them. To ensure that the transfer proceeds at the maximum rate, it is critical that ACKs are promptly sent back to the sender.

When the path that the ACK packets must transit is saturated by other connections (e.g. steady non-related upload traffic), TCP throughput will get delayed by default.

4.2.2.3.2 TCP SYNchronization

The SYN packet defined within the TCP protocol is used to set-up the connection before data transmission can take place. Packets meeting these criteria can be prioritized to enhance the performance of the network.

Figure VI illustrates the TCP header defined within this standard, with emphasis on the SYN flag within the header structure. The SYN flag is shown below in RED.

```
.-------------------------------+-------------------------------.
|          Source Port          |       Destination Port        |
|-------------------------------+-------------------------------|
|                        Sequence Number                        |
|---------------------------------------------------------------|
|                    Acknowledgment Number                      |
|-------------------+-+-+-+-+-+-+-------------------------------|
|  Data |           |U|A|P|R|S|F|                               |
| Offset| Reserved  |R|C|S|S|Y|I|            Window             |
|       |           |G|K|H|T|N|N|                               |
|-------+-----------+-+-+-+-+-+-+-------------------------------|
|           Checksum            |         Urgent Pointer        |
`---------------------------------------------------------------'
```

Figure VI – TCP header, including the SYN flag position in RED

The TCP SYN Flag has two possible values as follows:

Setting	*Meaning*	*Function*
0	Not Set	Packet is not a TCP SYN
1	Set	Packet is a TCP SYN

By observing traffic through the system and prioritizing of any packets with the TCP SYN Flag set, the TCP throughput can be enhanced. Allowing SYN packets to transit with priority will allow the associated traffic streams to transition to the data passing phase rapidly.

4.2.2.3.3 Layer 4 Classification

The figure below illustrates the classification process in taking in the default traffic stream and reordering the packets according to three example queues, TCP ACK, TCP SYN and the default best effort queue.

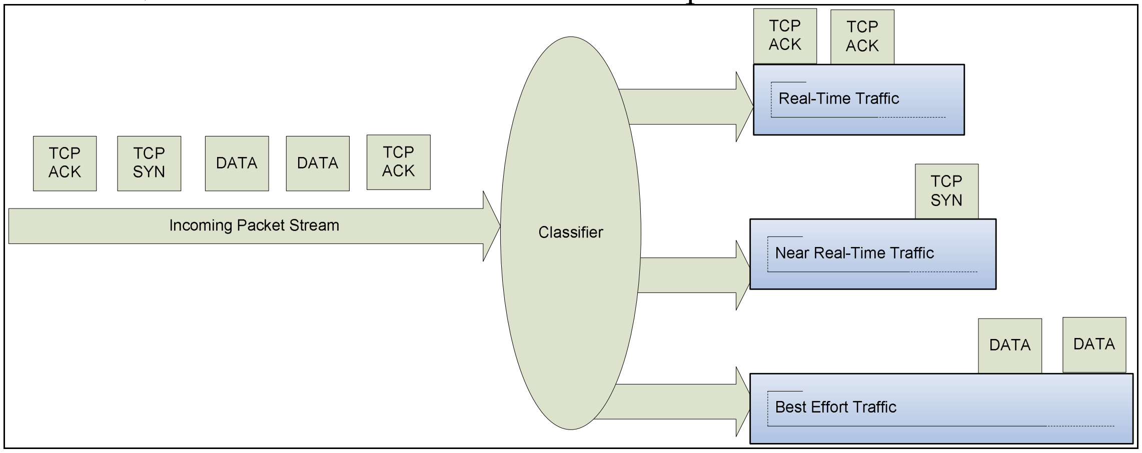

Figure VII – Layer 4 Classification

4.2.3 Example Flow classification

Applying the classification ideas presented in the previous sections can be accomplished via the following example rule-sets:

Traffic Type	***Characteristics***	***Priority***	***Example Rule***
VoIP	Needs real-time bi-directional service of 110kbps	HIGHEST	- Look for port numbers - Signaling = UDP 5060-5062 - Bearer = UDP 49150-
SSHV2/Telnet	Needs near real-time service of 9.6kbps	HIGH	- Look for port numbers - TCP port 22 and 23
FTP	Provide near real-time service	MEDIUM	- Look for port numbers - TCP port 20 and 21 - Prioritizing TCP ACK and SYN also assists this

Traffic Type	Characteristics	Priority	Example Rule
			protocol
HTTP	Provide best-effort service to this application.	LOW	- Look for port numbers - TCP port 80 - Prioritizing TCP ACK and SYN would assist this protocol indirectly
Best Effort	Provide best-effort service to this application.	LOW	- Flows not matching the signatures above fall into this category by default.
Undesirable Traffic	Provide less than best-effort, aka. low-effort, service to applications fitting this signature	LOWEST	-Traffic pattern/signature as per technical bulletins on Denial of Service, Virus alerts - Traffic patterns as per protocols in violation of the local acceptable use policy (AUP)

4.3 Station-Based Model

Bandwidth management can also be achieved with a station based model, which prioritizes traffic based on a station's class of service basis regardless of the presence of flows active on this or other stations in the network.

Traffic control utilizing this model ensures that quality of service is provided based upon the class-of-service (CoS) agreement that an end-user/subscriber currently holds with the service provider. This may be in the form of a Service Level Agreement (SLA) or less formal, for example: a Gold, Silver, & Bronze type classification.

Users with equal class-of-service parameters will obtain an equal share of network resources. In the presence of congestion, traffic for users with a higher class of service will still have all forms of their traffic transit the network, whereas traffic to/from lower priority users will not.

This model is far from ideal in today's converged infrastructures, where a given station on the network simultaneously supports numerous functions, and accordingly various protocols.

4.3.1 Priority Based Upon Station Address

To implement a station-based model, a mechanism must be developed to watch for traffic matching the address of the station(s) to be prioritized within the network. Typically, this field would be the source address or the destination address of the transmission, depending on whether the traffic is inbound or outbound with respect to the local network segment.

Priority may be given to network traffic received, sent or both on a given network station, or group of stations regardless of traffic/application type. Such stations may have users that require priority based upon their position in the organization, or be running applications that require a high quality of service link across the network. They may also be using applications which require highest/priority access across the network to support core business functions.

4.3.2 Example Station Classification

Classifiers could be created with the following rule-sets to manage network bandwidth based upon a station-by-station model.

Traffic Type	***Characteristics***	***Example Rule***
From Mgmt Station	Provide real-time bi-directional service to these stations at all times	Match on source IP of management station
From Engineering Department to subcontractor	Provide near real-time bi-directional service to these stations at all times	Match on source and destination addresses
To Sales Department	Provide near real-time inbound service to these stations at all times	Match on destination addresses
All Others	Provide best-effort access to this application	No match required, this is the default case

The figure below illustrates the classification process in taking in the default traffic stream and reordering the packets according to 3 example queues, Mgmt, Sales/Eng and Other.

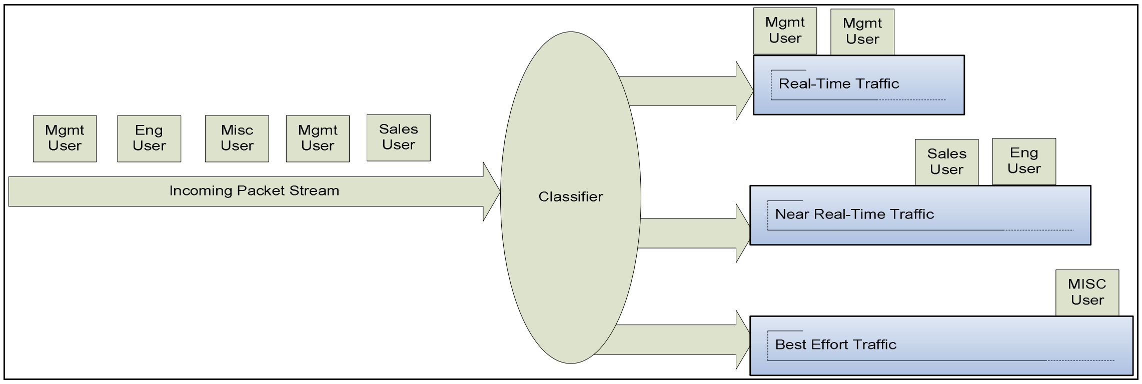

Figure VIII – Station Classification

4.4 Hybrid Model

This model consists of rule-sets within the bandwidth management device that can prioritize traffic based on a station basis, as well as, on a per-flow basis within a class of service.

This method attempts to balance the benefits of per flow prioritization with that of a station based priority implementation. The inherent challenge however is deciding which rule (flow or station based) should have ultimate priority.

4.4.1 Priority based on both flow and station address

In the most simplified case, users and applications can both categorized into high, medium and low priorities as in the following table:

User Priority vs. User	***Application Priority Vs. Application***
High – CEO	High – VoIP
Medium – Supervisor	Medium – Data Traffic
Low – Data Entry	Low – ICMP

This yields a total of 6 different classifiers for traffic, some based upon station address, while others are based upon application flow signature. The key decision to be made at this point is then whether the network should adopt a priority system where user priority is more important then application priority OR Application priority is deemed paramount over User priority.

In a configuration where user priority takes precedence over application priority, users that have low priority can find it difficult to effectively utilize their high priority applications. This is of course an undesirable effect.

Using the example table above, the following traffic priority mapping would be observed:

Priority	***Traffic Source***
1	CEO
2	Supervisor
3	Data Entry

In a configuration where application priority takes precedence over user priority, applications that have high priority will function well. However, this can impact the quality of service received at the high priority end points. An example would be the low priority that ICMP traffic from the CEO's workstation would receive.

Using the example table above, the following traffic priority mapping would be observed:

Priority	***Traffic Type***
1	VoIP
2	Data
3	ICMP

In a configuration where application priority and user priority are combined, the following traffic priority mapping would be observed:

Priority	***Traffic Source***
1	CEO or VoIP
2	Supervisor or Data
3	Data Entry or ICMP

This model will allow the network to function more effectively than with the default network configuration (e.g. without bandwidth management of any sort). Under congestion conditions, the implementation is less than ideal, since it allows at least some of the real-time traffic to queue behind arbitrary traffic based solely on station address.

5.0 Techniques for Bandwidth Management

There are varied solution types available for bandwidth management. These can be categorized as follows:

I. Bandwidth Management via hardware device
II. Bandwidth Management via integrated feature within a network device
III. Bandwidth Management via open-source software
IV. Bandwidth Management via commercial software

5.1 Bandwidth Management via Hardware

Network suppliers have several bandwidth options available to them, they may require their current vendors to supply products with integrated solutions or they may add in additional devices solely dedicated to bandwidth management. The merits of each solution will be discussed as follows.

5.1.1 Dedicated Hardware devices

This hardware generally consists of a device, with multiple network ports that match the incumbent network equipment, which is placed in-line with the network connection as shown in the figure below.

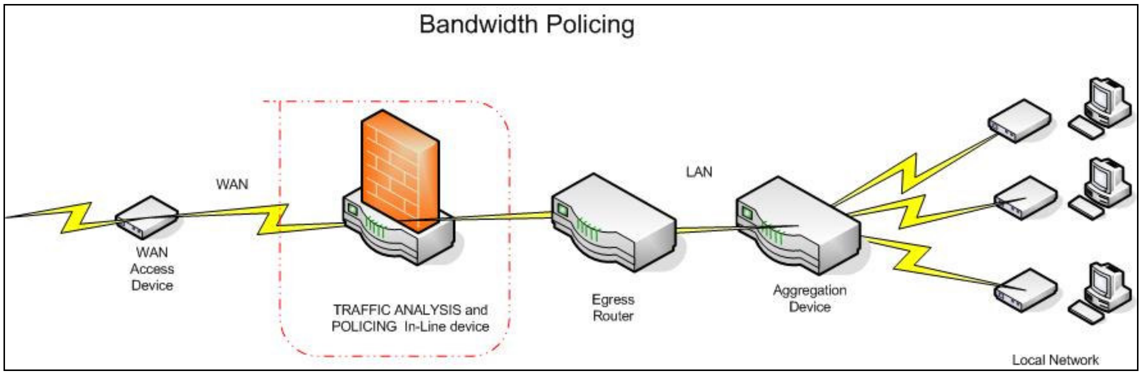

Figure IX – Dedicated Bandwidth Management Device Implementation

Once enabled and configured, the device inspects all traffic passing through it and enforces the configured policies to allow for the desired service levels. These levels may be prioritization of selected traffic, and/or de-prioritization of other traffic types. The inspection of traffic can be in either or both directions, making the versatility of the approach very flexible.

A sample selection of devices that reside in the category of inline dedicated bandwidth management device are:

I. IP Service Control system by Ellacoya Networks®[38]

II. Packetshaper product from Packeteer®[39]
III. Netenforcer by Allot[40]
IV. Policy Traffic switch from Sandvine[41]

5.1.1.1 IP Service Control System

This device provides the ability to manage bandwidth through the system and also generate associated reports on network usage on a per user and/or per application basis. This reporting functionality is useful as a feedback mechanism to capture network usage and subsequently apply them in bandwidth management policy decision making process.

5.1.1.2 Packetshaper®

The Packetshaper®[42] device is made by Packeteer© Inc. The devices is available in a number of different sizes/configurations from the Packetshaper® 1200, which would be applicable to 2Mbps networks, and the Packetshaper® 2500, which would service the target audience with its 10Mbps capacity. This device provides for efficiency calculations and also automated alerts based upon traffic signature observed within the network.

5.1.1.3 Netenforcer®

The Netenforcer® device is made by Allot Inc This device has been designed for small/medium sized networks. It utilizes deep packet inspection techniques to monitor and control bandwidth within the network. The device applicable to the target audience would be the AC-402, which will support the 10Mbps target bandwidth.

5.1.1.4 Policy Traffic Switch

This device is made by Sandvine® Inc. The Policy Traffic Switch 8210, is an inline device that has been designed with extensive CPU and memory resources onboard. In addition, it contains proprietary algorithms to collect, analyze, and assist with the interpretation of network traffic on a per-flow basis.

The device performs stateful-flow inspection on all flows observed through the device. This implementation allows for reconstruction of the complete (bidirectional) flow pattern for all network traffic scheme.

5.2 Integrated Bandwidth management

This implementation consists of a network device that was already resident within the network to become enabled for the bandwidth management task in addition to their usual functions.

There is a dependency here on whether a network device supports this type of advanced feature. Assuming a device supports this functionality, there is also another consideration that should be carefully weighed. That is the performance of the device may be significantly impacted by the activation of this additional feature. Most commonly, this means an increase in CPU utilization, a key metric that operators strive to keep as low as possible.

Sample selections of implementations that reside in this category are:

I. Network Based Application Recognition (NBAR)[43] by Cisco Systems™

II. Network Flow (NETFLOW) by Cisco Systems™ [44]

5.2.1 NBAR®

The Network Based Application Recognition (NBAR®) feature is provided within select Cisco Internetworking Operating System(IOS) versions. This proprietary feature can have a significant impact upon the Central Processor Unit (CPU) performance of any Cisco™ device when enabled. Use of this feature also drives the network operator to obtain and maintain NBAR® aware Cisco™ devices throughout the network.

5.2.2 NetFLOW®

The Cisco IOS NetFlow® is a similar feature that provides the ability to track information about network users and applications, peak usage times, and traffic routing. This feature is proprietary to Cisco™ and is utilized by applications, open-source and otherwise, as a foundation to implement traffic accounting and subsequent shaping based upon the observed traffic in the network.

Policies can be created and implemented based upon the information that NetFLOW® provides, however this data is only available from Cisco™ devices, and thus there is a dependency on vendor type within the network in order to utilize this function.

5.3 Open Source Bandwidth management

This implementation consists of a personal computer(PC) hosting an open source operating system, most commonly a Linux™ variant, and running an open source application within this operating system to manage bandwidth for the network.

The bandwidth management device would have dual network interface cards (NICs) installed and configured, to provide for in-line placement within the network, similar to the dedicated hardware solution explored above, which also have an 'in' and 'out' port.

The first of these NIC cards would interface with the local network(LAN), and the other would form the connection to the Wide Area Network (WAN). This allows for traffic to be evaluated and managed as it flows between the local loop and the WAN connection.

The main differentiator of this solution and the dedicated device approach discussed above would be in the area of capital and operations costs. Being open-source in nature the cost, will be significantly lower than the dedicated, and largely proprietary, hardware approaches discussed above in terms of capital cost (CapEx). The open-source nature also avoids the recurring software maintenance fees (OpEx) inherent in the major of the commercial solutions in the marketplace today.

Sample selections of applications that reside in this category are:

I. Monowall[45]©
II. Dummynet[46]
III. Master Shaper[47]
IV. IPTables[48]

5.3.1 Monowall©

Monowall© is a project that was initiated to create a complete firewall implementation based upon FREE-BSD[49], with the intention of mirroring the features of a commercial firewall. This application has been developed for desktop type systems as well as the smaller embedded PC environment.

The application build is complete with a web server, which is utilized for GUI management of the system. There is also an option to install the application to a CD ROM and then utilize this CD ROM in the booting of the system. This removes the requirement to have a machine (re)built under the LINUX™ operating system.

5.3.2 Dummynet

Dummynet works in-conjunction with IP Firewall (IPFW) rules. The implementation uses two main concepts: queues and pipes. These components are designed to simulate the effects of network bandwidth limitations, propagation delays, and packet loss.

Queues represent instead queues of packets, associated with a weight, which share the bandwidth of the pipe they are connected to proportionally to their weight.

Pipes are fixed-bandwidth channels. Each pipe and queue can be configured separately, so you can apply different limitations/delays to different traffic according to the IPFW rules (e.g. selecting on protocols, addresses and ports ranges, interfaces, etc.).

Pipes and queues can be created dynamically, so using a single set of rules you can apply independent limitations to all hosts in a subnet, or to all types of traffic, etc. You can also configure the system to build cascades of pipes, so you can simulate networks with multiple links and paths between source(s) and destination(s).

5.3.3 Master Shaper

Master Shaper is a network traffic shaper application under Linux™, which provides a GUI Web Interface with which a network operator can implement Quality of Service (QoS) functions.

It allows users to use traffic shaping mechanisms using a graphical user interface (GUI) and via the definition of bandwidth pipes and filters. This application can also output data in graphical format to represent current bandwidth usage and distribution.

5.3.4 IPTables

IPTables is a command line program used to configure the standard Linux™ packet filtering rule set.

5.4 Non-Open Source Bandwidth management

This implementation consists of a computer hosting windows and running a variant of a non-open source application to manage bandwidth. The device would have dual network interface cards (NICs), to provide for in-line placement within the network, similar to the dedicated hardware solution discussed above. This allows for traffic to be evaluated and shaped as it flows between the local loop and the WAN connection.

Sample selections of applications that reside in this category are:

I. Traffic Controller[50]
II. Softperfect™ Bandwidth Shaper[51]
III. DU Super Controller[52]
IV. Inetshaper[53]

5.4.1 Traffic Controller

This is a windows based application, which is available to support networks of the size intended for analysis within this paper. The implementation provides for prioritization of UDP and/or TCP traffic. The graphical user interface is user-friendly for the operator to create rules and verify configuration.

Approximate Pricing: $189.95

5.4.2 SoftPerfect™ Bandwidth Shaper

SoftPerfect™ Bandwidth Shaper is a traffic management tool running under the Windows operating system. It offers bandwidth control and quality of service capabilities based on built-in prioritized rule sets. These rules can specify a bandwidth and throughput limits to specified IP addresses, and also by port classifications.

Approximate Pricing: $99

5.4.3 DU Super Controller

Control over download and upload throughputs, and ability to limit throughputs in either direction. Ability to prioritize traffic flow to maximize throughput on asymmetric links. Also has feature set to disable rules for 'local' traffic, and instead pass this without modification to other stations on the local network.

Approximate Pricing: $75

5.4.4 Inetshaper

The Inetshaper application allows for traffic shaping, filtering capabilities, access accounting, as well as, traffic statistics presentation. Classification of users can be accomplished via IP address, as well as, MAC address. This application allows for the prioritization of traffic among users on the network.
Approximate Pricing: $75

6.0 Bandwidth Management Implementation

Of the hardware and software based solutions defined in section 5, there are many trade-offs that must be taken into account to provide the most applicable solution for the target audience.

6.1 Target Network Requirements

First a target network must be selected, and for the purposes of this paper, that audience has been previously identified as meeting the characteristics of a small ISP. The characteristics of these networks can be categorized as follows:

6.1.1 Network throughput requirements

The network bandwidth for such networks is assumed to be in the megabits range, however generally characterized as below 10Mbps in any one direction. Many operate at significantly lower speeds traceable to the legacy "T1" style service that was common for WAN links in the recent past. In those cases, the maximum rate possible would be 1.544Mbps in any one direction.

6.1.2 Network station requirements

The network population on such networks can range from a dozen, or less users, on into the 100's range. Surveying several small live ISP operations in Alberta and Ontario has yielded an average of 200-250 subscribers on these networks.

6.1.3 Application support requirements

The application support requirements of the small ISP are diverse in nature and can range from traffic that requires real-time transport to that which can operate well with only best-effort transport available.

Typically, the small ISP themselves may not add or directly provide the type of advanced services/protocols that require specialized quality of service. Instead they tend to concentrate on a pure access, or 'pipe', service.

That being stated, there remains a vested interest to provide appropriate QoS to these advanced applications as they transit the local network in order to build and retain the customer base. Thus a provider that deploys a bandwidth management solution will benefit on an ongoing basis in terms of customer satisfaction.

6.2 Applicability of the Presented Solutions

Each of the bandwidth management solutions presented in this paper will accommodate the expected network throughput of 10Mbps or less in each direction. The hardware based solutions, both in the form of dedicated in-line devices, as well as, with vendor-specific integrated features, tend to be designed to accommodate networks of much greater throughputs in the multi-gigabit range.

The target audience can accordingly be accommodated with the solutions that are suitable for lower bandwidth networks. These are the open-source and non open-source implementations.

All of the bandwidth management solutions described in this paper will accommodate the number of stations resident on the target audience's networks. The hardware based solutions, both in the form of dedicated in-line devices, as well as, with vendor-specific integrated features tend to be designed to accommodate networks of much greater populations, 1,000's to 10,000 users. The target audience will provide network service to a few hundred, or in many cases, fewer than 100 users simultaneously on a single network segment.

Each of the bandwidth management solutions presented in this paper will accommodate the application support required by the target audience. All solutions can provide both prioritization of important traffic and de-prioritization of nuisance, or unimportant, traffic to varying degrees.

While, the hardware based solutions do present more comprehensive feature-sets in implementing their solutions, the feature-sets offered by the open source and non open-source software implementations are already sufficient for meeting the requirements of most small operators.

The smallest of the operators in the target audience will also tend to prefer the simplified approach that the software implementation provides, as their depth of network skills is typically very limited in nature.

For the networks under consideration within this paper, the software based approaches using non open-source implementations present a tangible, yet feasible cost to implement. The open-source implementation however is more suitable due to its lower initial cost, as well as, its lower ongoing cost in terms of support and upgrades to functionality. This style of implementation will provide for an increase in network performance while presenting very low

CapEx and OpEx commitments during the implementation and operation phases.

7.0 Proposed Implementation

In proposing a solution for implementation on the target audience's networks, the following additional factors have been considered:

I. portability between network architectures and equipment vendors
II. overall cost of each bandwidth management strategy
III. performance cost/benefit ratio of each strategy

7.1 Target Audience

While the focus of this paper is on the application of a bandwidth management solution for the Small ISP, there are several similar small-medium scale networks that can achieve significant benefit from the bandwidth management techniques described as well. A sample of these network types is presented in the following sections:

7.1.1 Small ISP Operators

These networks purchase their bandwidth from a larger, and sometimes competing, provider to re-sell that bandwidth to their subscriber base. Minimizing the size of the pipe that they need to obtain from these larger providers has monetary advantages on a recurring basis. These recurring savings may be utilized to defer the OpEx and/or CapEx of the selected solution.

Small operators will lease circuits/connections from the larger network operators in the geographic area. In Alberta providers that would supply small internet providers would include Telus, Sprint, AllStream and Shaw's Big Pipe[54].

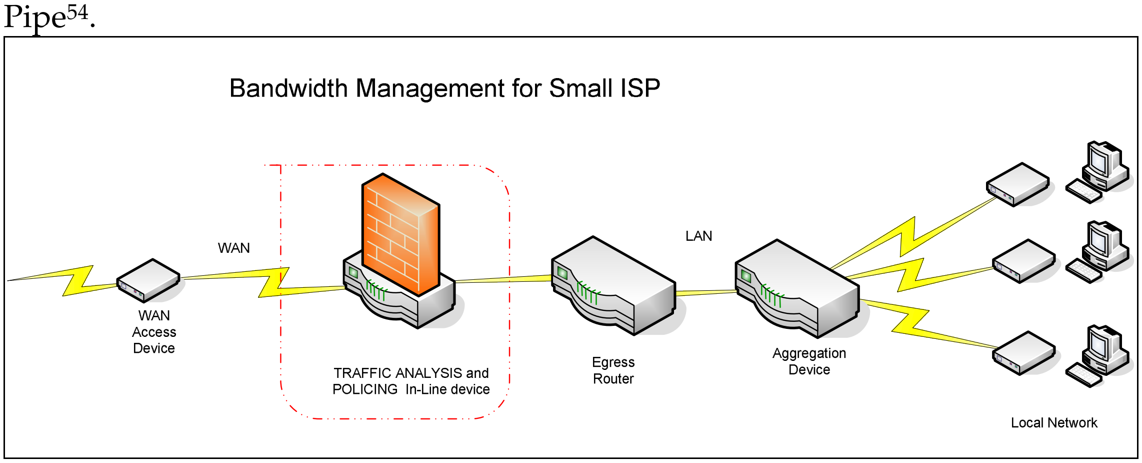
Bandwidth Management for Small ISP
WAN
WAN Access Device
TRAFFIC ANALYSIS and POLICING In-Line device
Egress Router
LAN
Aggregation Device
Local Network

Figure X – Bandwidth Management for Small ISP Applications

7.1.2 Temporary industrial installations

These networks typically have a Wide Area Network (WAN) connection over microwave or SATCOM links. The cost on a per-byte basis of this style of connection is very high. Controlling the traffic that must transit this link provides for an ongoing opportunity for cost savings.

Quality of Service and maximum efficiency of the WAN connection would also be applicable to remote locations. Within Alberta's industrial economy, these applications would include remote production locations, the geographical placement of which may necessitate connection by relatively low speed, and costly links. Given the transient nature of the majority of these work sites, links across a medium such as satellite would be more favorable than a hard-wired infrastructure.

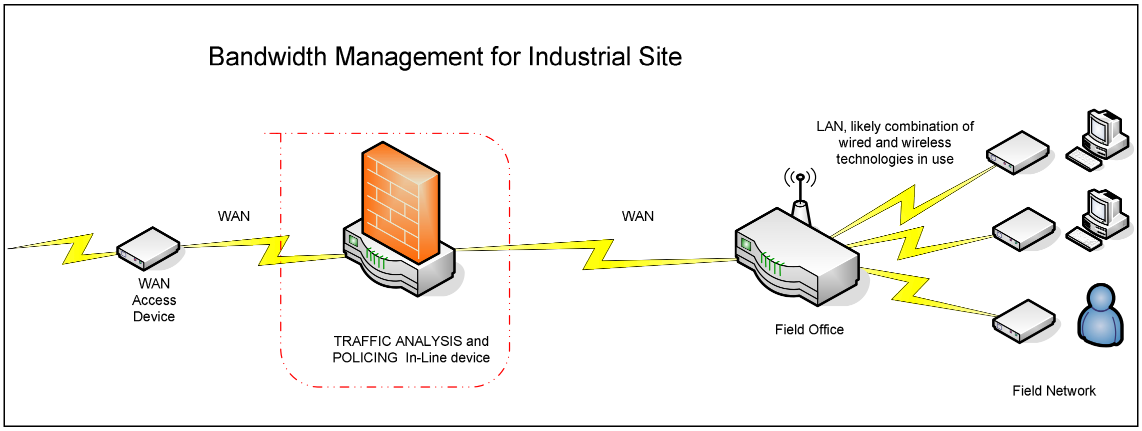
Bandwidth Management for Industrial Site
WAN
WAN
WAN Access Device
TRAFFIC ANALYSIS and POLICING In-Line device
LAN, likely combination of wired and wireless technologies in use
Field Office
Field Network

Figure XI – Bandwidth Management for Industrial Applications

7.1.3 Small Educational Institutions

These networks typically have a WAN connection over fiber optics or broadband connection which may connect to the public internet, or to other educational institutions. In the latter cases, the costs of using the link will not be the primary concern, however optimizing the use of the inter-institutional links is of greater importance. A ubiquitous example of this type of network structure would be public libraries across the province.

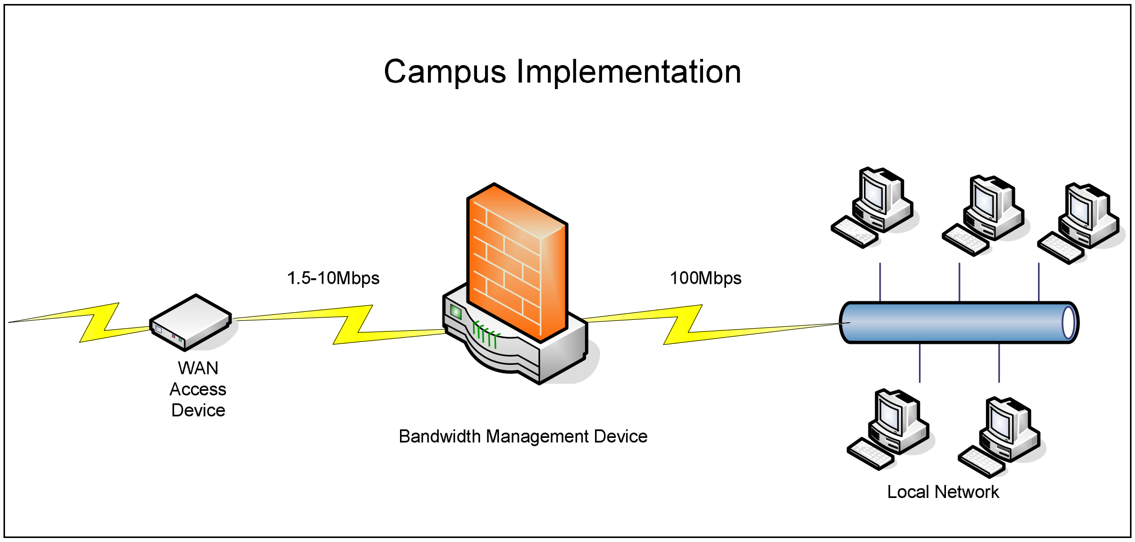

Figure XII – Bandwidth Management for Campus Applications

7.2 Parameters Considered

There are several metrics that can be utilized in the evaluation of the available solutions. These metrics are both technical and non-technical in nature. The design goals of the solution described in this paper are to implement a system that is low cost, low complexity and provides longevity for the network provider.

7.2.1 Network Performance Gain

The implementation of the bandwidth management function within the network will allow traffic flows to be controlled and quality of service to be implemented as desired by the Small ISP operator.

By optimizing traffic within the network, the operator will observe an increase in customer satisfaction, increased customer retention and avoid expensive pre-mature investment in additional capacity.

The design objective is to propose a solution that would provide a significant overall network performance gain to the Small ISP operator. The implementation in section 7.3 achieves this goal.

7.2.2 Portability of Solution

The ideal solution would require minimal rework or supplementary cost when changes in the network take place. Example of modifications within the network would be the upgrade of existing network equipment, the changing of platforms within the current vendor, implementation of a new vendor's solution, and implementation of a future technology.

While the Small ISP network is not likely to see as many changes as larger networks, factors such as: equipment obsolescence, equipment failure and eventually capacity constraints will inevitably drive modifications to the network over time.

The design objective is to propose a solution that would require minimal, if any, reconfiguration in these events. The implementation in section 7.3 achieves this goal.

7.2.3 Cost of Solution

There are two main components to consider in examining the cost of the solutions that follow. These are capital expenditure (CapEx) and also operating expenditure (OpEx).

CapEx costs are borne upon initial procurement and deployment of the solution. These costs are directly related to the procurement of software, hardware and associated materials. In addition, significant costs may be incurred initially in the installation and commissioning of the components identified above.

OpEx costs are borne on an ongoing basis through the life cycle of the solution within the network. These costs are related to operating the solution on an ongoing basis. A large component of these recurring costs would be in the form of software licensing fees that the major network equipment manufacturers charge on a yearly basis. For commercial networking solutions, the annual cost of these fees is typically in the range of 8-10% of the CapEx cost.

The design objective is to propose a solution that would require minimal CapEx, as well as, OpEx. The implementation in section 7.3 achieves this goal.

7.2.4 Complexity of Solution

The networking expertise resident within a Small ISP's core staff will vary from operator to operator. Typically, this expertise is found on a much more limited basis than within the organizations of larger operators. A solution that is straightforward to implement, with minimal operating system interaction, will be preferred by the target audience.

The design objective is to propose a solution that would require minimal networking expertise to implement. The implementation in section 7.3 achieves this goal.

7.3 Proposed Implementation

Applying the considerations outlined in sections 6 and 7.2 to the network of the target audience and analyzing the various hardware and software options presented in the previous sections, an open source implementation is proposed.

Of the open source implementations examined, the Monowall open-source application has been selected and trialed based upon its feature-set, ease of set-up, portability into the future and user-friendly interface. This application is an Open source package and is available to use within Linux™, OpenBSD or direct-boot from CD-ROM.

The other components within the solution consist of a single commercial off-the-shelf (COTS) PC running the open source operating system Linux™. Upon this Linux™ device, there are two Network Interface Cards (NICs) installed. These provide for an input and output network interface at a line rate of 100Mbps, and allows for the device to typically be placed at the egress point of the network, inline with the network.

The bandwidth management device would be installed at the point of demarcation between the user's network and the Wide Area Network (WAN). Placement of the device at the egress point in the network minimizes replication required to support multiple devices within various locations within the local network, thus reducing the costs of implementation.

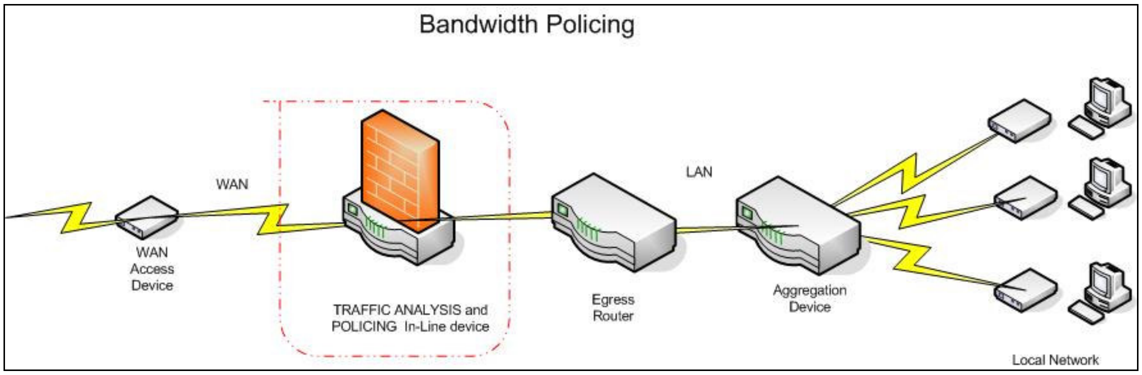

Figure XIII – Placement of the Bandwidth management device

Placement of the device at this point also is important in that it removes any dependency upon the physical network in use. The solution is equally applicable for networks that run xDSL, Cable, Wireless or fixed wireless implementations.

Each of these networks may run proprietary protocols on their local links, however in the end translate all traffic to standard IPV4[55] packets for transmission to the WAN connection. It is at this point that the solution may be introduced and be able to leverage this position to observe packet flow in each direction as it enters and leaves the network.

The station is configured to bridge, which allows packets to enter one NIC, be evaluated by the traffic control software resident upon the station, and in turn transit to the outbound interface based upon a pre-defined set of rules and policies[56].

7.3.1 Bill of Materials

To summarize a bill of materials, the solution would call for the following components:

I. Dedicated PC of at least 1GHz CPU, with CD-ROM drive
II. Dual Network Interface Cards (NICs) installed
III. Monowall application installed on this PC

These components are further illustrated in the following diagram.

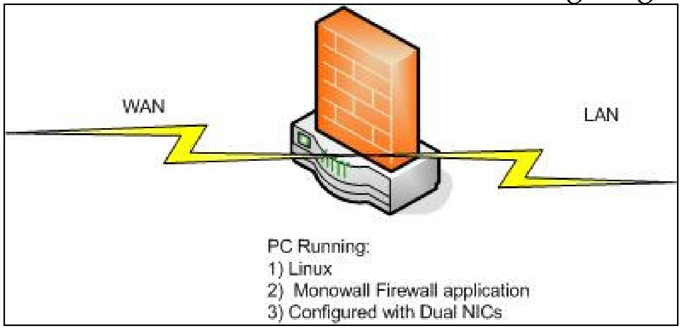
WAN
LAN
PC Running:
1) Linux
2) Monowall Firewall application
3) Configured with Dual NICs

Figure XIV – Placement of the Bandwidth management device

7.3.2 Supporting Applications

The following applications are recommended for use in the development of the prototype implementation and subsequent test phases:

Open-Source/freeware

I. IPERF[57] – Provides for a server—client connection in order to transport UDP or TCP test packets with a variety of settings for throughput, packet size, characteristics.

II. FPING[58] – Provides for an ICMP generating utility that can be configured to send traffic with varying sizes, delays and characteristics useful for testing and characterization purposes.

III. NETPERF[59] – Similar in functionality to IPERF(above). Provides for a server—client connection in order to transport UDP or TCP test packets with a variety of settings for throughput, packet size, characteristics

IV. ETHEREAL[60] – Provides for packet capture and analysis functionality. Also can be utilized to reconstruct packet streams such as the RTP streams inherent in VoIP conversations.

Commercial

I. COMMVIEW[61] – Provides similar functionality to ethereal, with some expansion in the area of traffic accounting, which may be useful in the test phases of this project.

II. COMMTRAFFIC[62] – Provides functionality to complement commview or ethereal in the accounting of traffic as it transits the network.

7.3.3 Recommended Configuration

For the implementation phase, the configuration recommended would require a minimum of 4 service classes as follows :

I. Real-time Traffic
II. Near Real-time Traffic
III. Best Effort Traffic
IV. Nuisance Traffic

Additional traffic classes, to further classify traffic beyond the bounds proposed above can certainly be supported through means of additional configuration within the bandwidth management application.

The real-time traffic class would be utilized to support applications that run over protocols requiring extremely low latency, fixed jitter and robust performance. Examples of this traffic type are VoIP calls and streaming Video sessions.

The near real-time traffic class would be utilized to support applications that run over protocols requiring relatively low latency & jitter. These protocols can handle more delay in the network than the real-time class can tolerate. Practical examples of traffic belonging in this class are SSH and Telnet sessions.

The best effort traffic class would be utilized to support traffic not matching any of the classifiers configured. The majority of traffic within this class will be bulk data transfer traffic, with a large tolerance to jitter, delay and even a moderate tolerance for packet loss impairments. Examples of this traffic type are HTTP browsing and Network News Transfer Protocol (NNTP).

The nuisance traffic class would be utilized to support the de-prioritization of traffic deemed undesirable. The rule-set relevant to this class of traffic may call for a throttling of resources available to the traffic, or an all-out blocking of the traffic depending on the negative impact that the traffic is deemed to have on the LAN and WAN network resources. Examples of this traffic type are virus generated probes and denial of service floods.

8.0 Project Extensions

Beyond the implementation activities proposed in section 7 above, there are several extensions that could be pursued based upon the architecture and solution presented within this paper to increase the capabilities of the solution.

8.1 Areas for future research

Some recommendations for future areas of research related to this project are:

I. Implementation of a reports function to provide historical data on usage patterns

II. Implementation of a packet 'mangler' function to remark TOS/DSCP

III. Analysis of a migration strategy of this solution to support an IPV6 environment from the current IPV4 implementation

8.1.1 Reports/Accounting function

Implementation of a reports function to provide a comprehensive real-time and historical view on usage patterns would expand the scope of the proposed solution to bring in a feedback mechanism that the network operator could leverage in order to adjust network policy on a regular basis. The accounting of these parameters could be on a per-flow basis, allowing the tracking of per-application and/or per-user.

Such information is available via the use of 3rd party applications today, however many of these are non-open source in nature. An analysis of the available applications could be undertaken for future work.

8.1.2 Packet mangling Function

Implementation of a packet 'mangler' function, to provide the ability to manipulate packets that transit the system. Specifically, the remarking of the type-of-service(TOS)/DSCP fields would be a complementary extension to the prioritization functions inherent to the solution proposed within this paper.

The same architecture could be leveraged to provide for a remarking function as well the core prioritization function to ensure traffic is marked according to local policy, rather than its original policies.

This presents potential benefits of better processing of this traffic once it is on the WAN connection, and also for traffic entering the network to ensure local

Quality of Service policies are maintained. The main fields to be modified would be those resident within the IP header of the TOS/DSCP flags.

8.1.3 IP Version 6 (IPV6) support

Analysis of a migration strategy that could be applied to the solution to support an IPV6[63] environment as a future migration path from the current IPV4 implementation could be undertaken. The prevalence of IPV6 networks will increase over the next few years, making such an analysis increasingly relevant.

There are two environments that will develop in most networks and suitability of the solution to migrate support for these could be analyzed in detail. These two environments can be described as:

1. dual IPV4/IPV6 environment
2. pure IPV6 environment

8.1.3.1 Dual IPV4/IPV6 environment

As network operators begin turning towards IPV6 solutions, they will have significant investment in IPV4 infrastructure. As operators introduce the first IPV6 components into their network, there will be a need for translation back to IPV4 to transit networks or segments of the network not yet converted to IPV6.

This dual environment is expected to exist for an extended period of time and so there is a requirement to study adaptations necessary to the proposed solution in order to allow it to co-exist in this upcoming environment.

8.1.3.2 Pure IPV6 environment

In "Greenfield" scenarios where new networks are put into place, or when a transition to IPV6 has been completed there will be a requirement to adapt the proposed solution to accommodate a purely IPV6 traffic flow. Thus the requirement in this case is to study how the proposed solution could migrate to a pure IPV6 environment.

9.0 Conclusions

The introduction of bandwidth management device and classification of traffic as discussed can significantly enhance the overall performance of the Small ISP operator’s network.

This has numerous benefits including: maximum efficiency for the network provider, maximal fair-sharing of available bandwidth and the ability to offer improved quality of service to specific users or applications on the network.

This will also result in non-technical benefits such as decreased WAN costs, decreased capital and operational expenditures, as well as, increased levels of customer satisfaction/retention.

As networks grow in the future with increased volume of converged traffic types, bandwidth management through prioritization will become an increasingly important component of the best-practices operation of the network.

The selection of a solution that fits the requirements of the target networks well, at minimal cost and with minimal front-end time configuration time for a given network. The straightforward nature of the Graphical User Interface inherent in the solution will aid in the configuration phase.

The implementation phase discussed in this paper will allow future MINT students to implement the strategies explored and gauge impact to network performance.

The recommendations for follow-on research will ensure the proposed solution is applicable well into the future and provides further value to the target audience.

Appendix A – Monowall User Interface

This appendix provides some sample configuration views from the Monowall graphical user interface (GUI)[64].

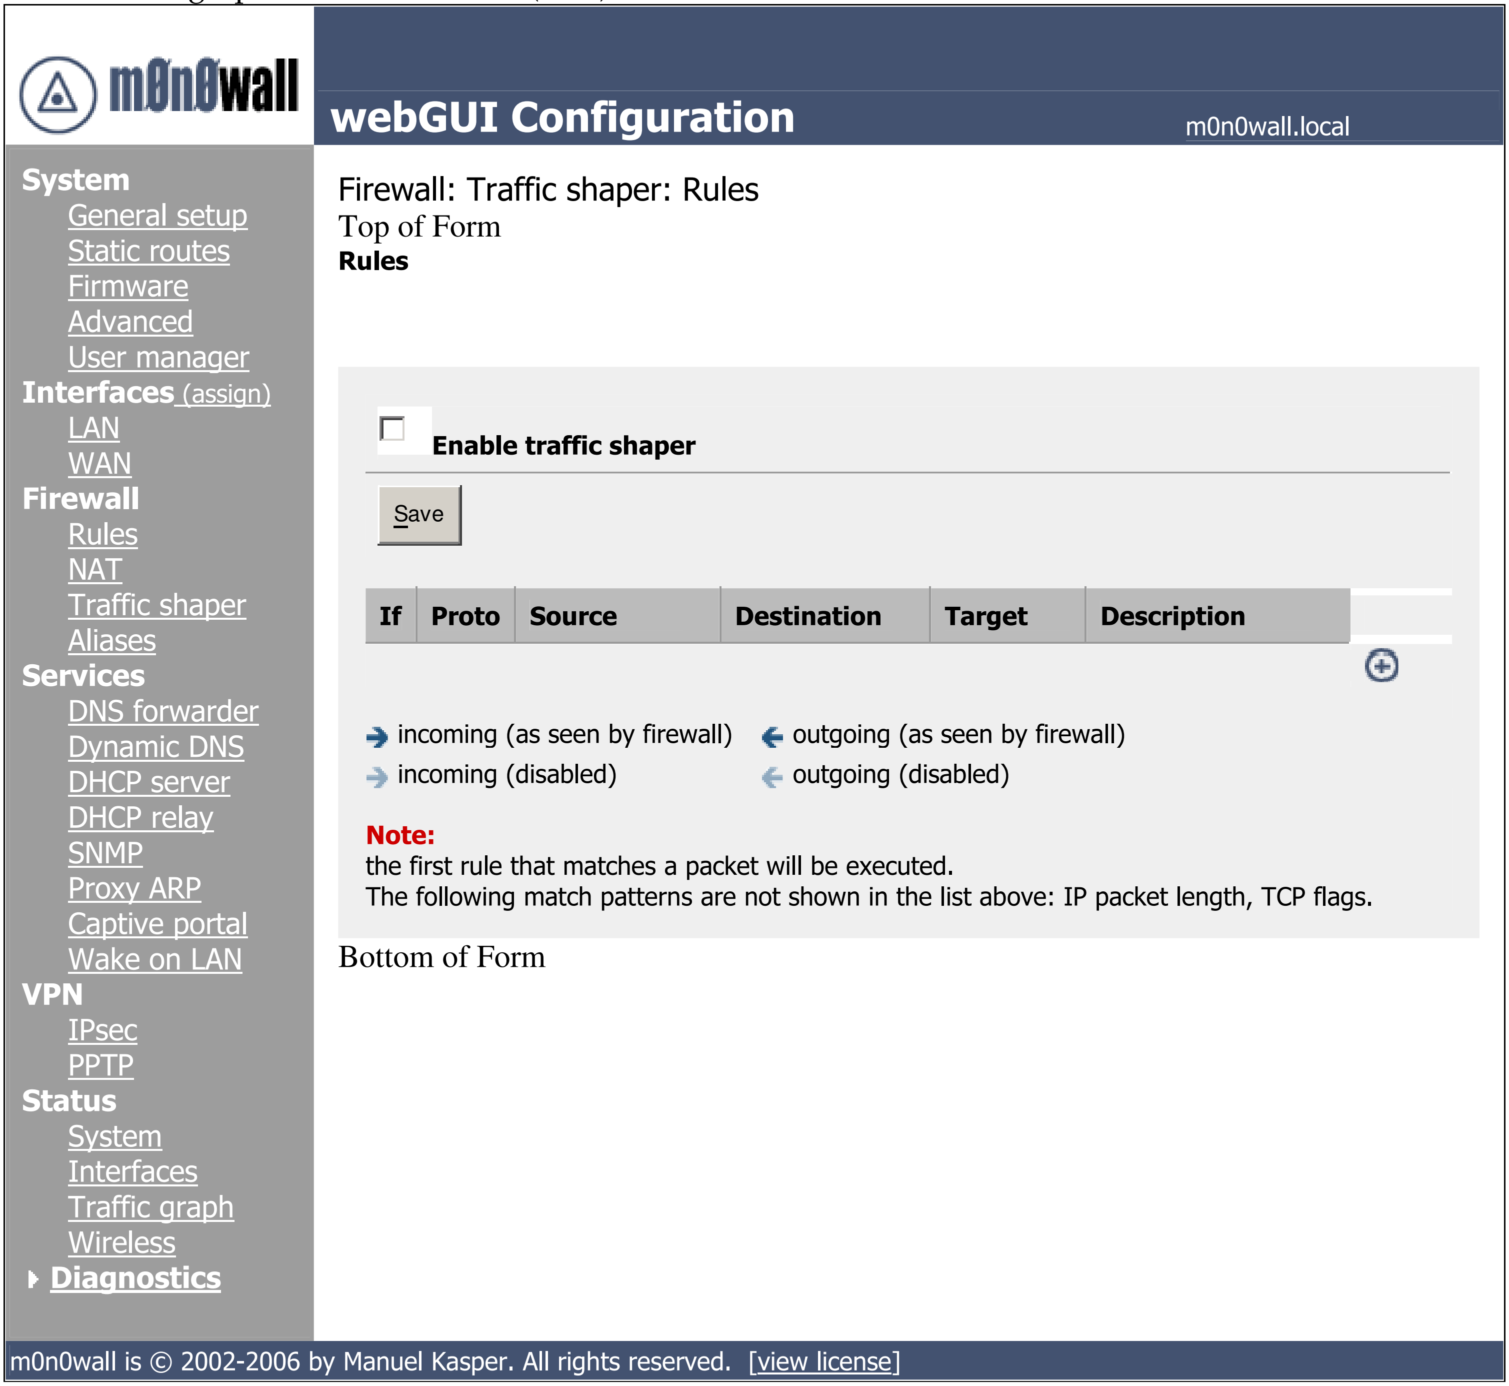

Firewall: Traffic shaper: Rules

Top of Form

Rules

□ **Enable traffic shaper**

If	Proto	Source	Destination	Target	Description

incoming (as seen by firewall) outgoing (as seen by firewall)

incoming (disabled) outgoing (disabled)

Note:
the first rule that matches a packet will be executed.
The following match patterns are not shown in the list above: IP packet length, TCP flags.

Bottom of Form

m0n0wall is © 2002-2006 by Manuel Kasper. All rights reserved. [view license]

Firewall Shaper Rules - Queues

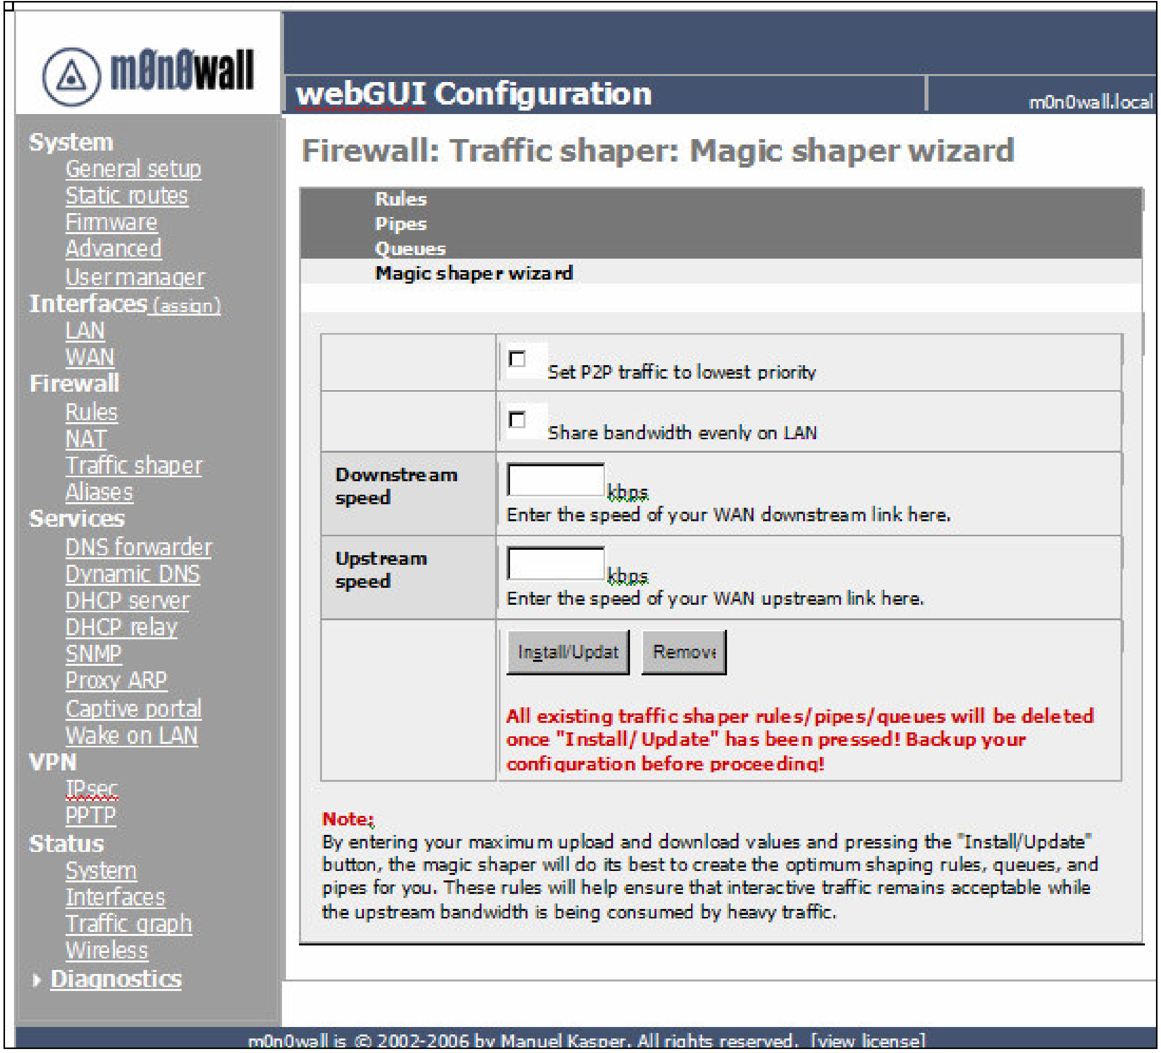

Firewall: Traffic shaper: Magic shaper wizard

	☐ Set P2P traffic to lowest priority
	☐ Share bandwidth evenly on LAN
Downstream speed	kbps Enter the speed of your WAN downstream link here.
Upstream speed	kbps Enter the speed of your WAN upstream link here.
	Install/Updat Remove All existing traffic shaper rules/pipes/queues will be deleted once "Install/Update" has been pressed! Backup your configuration before proceeding!

Note:
By entering your maximum upload and download values and pressing the "Install/Update" button, the magic shaper will do its best to create the optimum shaping rules, queues, and pipes for you. These rules will help ensure that interactive traffic remains acceptable while the upstream bandwidth is being consumed by heavy traffic.

m0n0wall is © 2002-2006 by Manuel Kasper. All rights reserved. [view license]

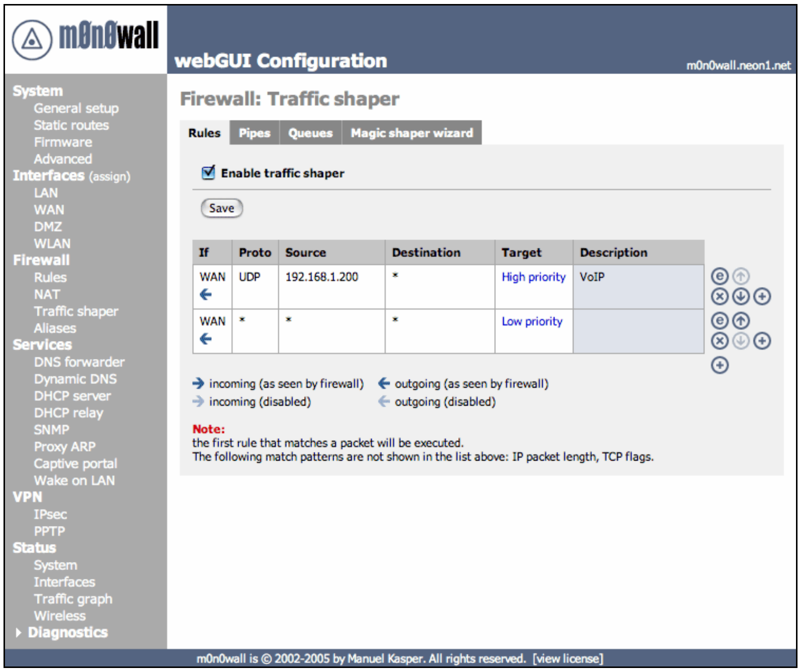

Firewall: Traffic shaper

☑ Enable traffic shaper

If	Proto	Source	Destination	Target	Description
WAN ←	UDP	192.168.1.200	*	High priority	VoIP
WAN ←	*	*	*	Low priority	

→ incoming (as seen by firewall) ← outgoing (as seen by firewall)
→ incoming (disabled) ← outgoing (disabled)

Note:
the first rule that matches a packet will be executed.
The following match patterns are not shown in the list above: IP packet length, TCP flags.

m0n0wall is © 2002-2005 by Manuel Kasper. All rights reserved. [view license]

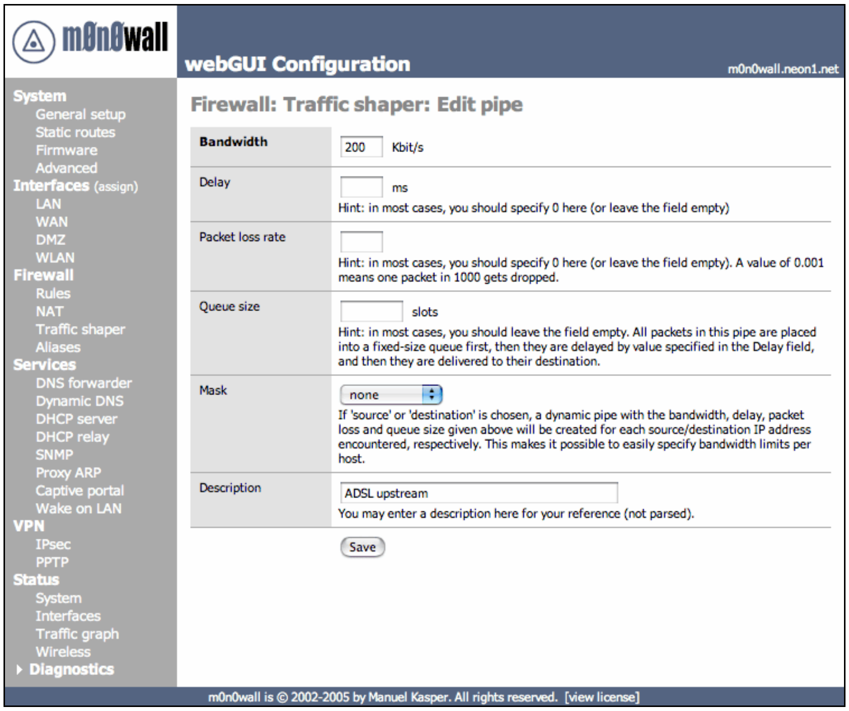

Firewall: Traffic shaper: Edit pipe

Bandwidth	200 Kbit/s
Delay	ms Hint: in most cases, you should specify 0 here (or leave the field empty)
Packet loss rate	Hint: in most cases, you should specify 0 here (or leave the field empty). A value of 0.001 means one packet in 1000 gets dropped.
Queue size	slots Hint: in most cases, you should leave the field empty. All packets in this pipe are placed into a fixed-size queue first, then they are delayed by value specified in the Delay field, and then they are delivered to their destination.
Mask	none If 'source' or 'destination' is chosen, a dynamic pipe with the bandwidth, delay, packet loss and queue size given above will be created for each source/destination IP address encountered, respectively. This makes it possible to easily specify bandwidth limits per host.
Description	ADSL upstream You may enter a description here for your reference (not parsed).

m0n0wall is © 2002-2005 by Manuel Kasper. All rights reserved. [view license]

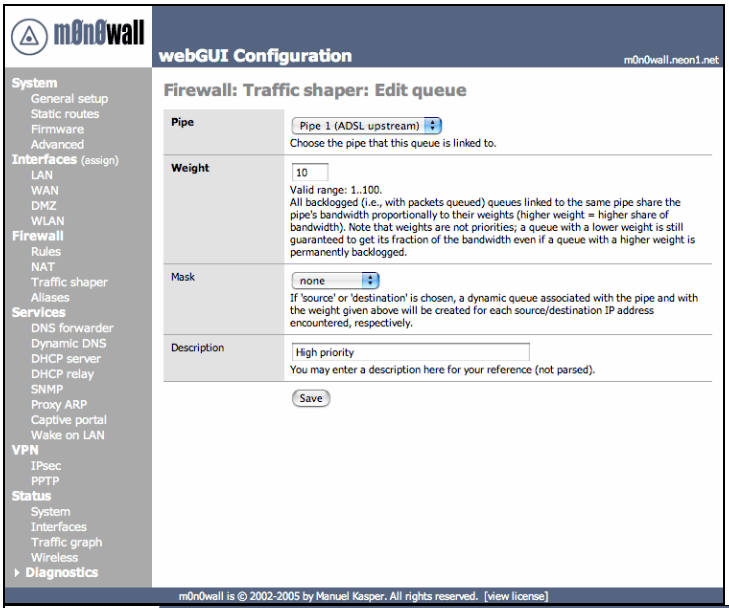

Firewall: Traffic shaper: Edit queue

Pipe	Pipe 1 (ADSL upstream) Choose the pipe that this queue is linked to.
Weight	10 Valid range: 1..100. All backlogged (i.e., with packets queued) queues linked to the same pipe share the pipe's bandwidth proportionally to their weights (higher weight = higher share of bandwidth). Note that weights are not priorities; a queue with a lower weight is still guaranteed to get its fraction of the bandwidth even if a queue with a higher weight is permanently backlogged.
Mask	none If 'source' or 'destination' is chosen, a dynamic queue associated with the pipe and with the weight given above will be created for each source/destination IP address encountered, respectively.
Description	High priority You may enter a description here for your reference (not parsed).

m0n0wall is © 2002-2005 by Manuel Kasper. All rights reserved. [view license]

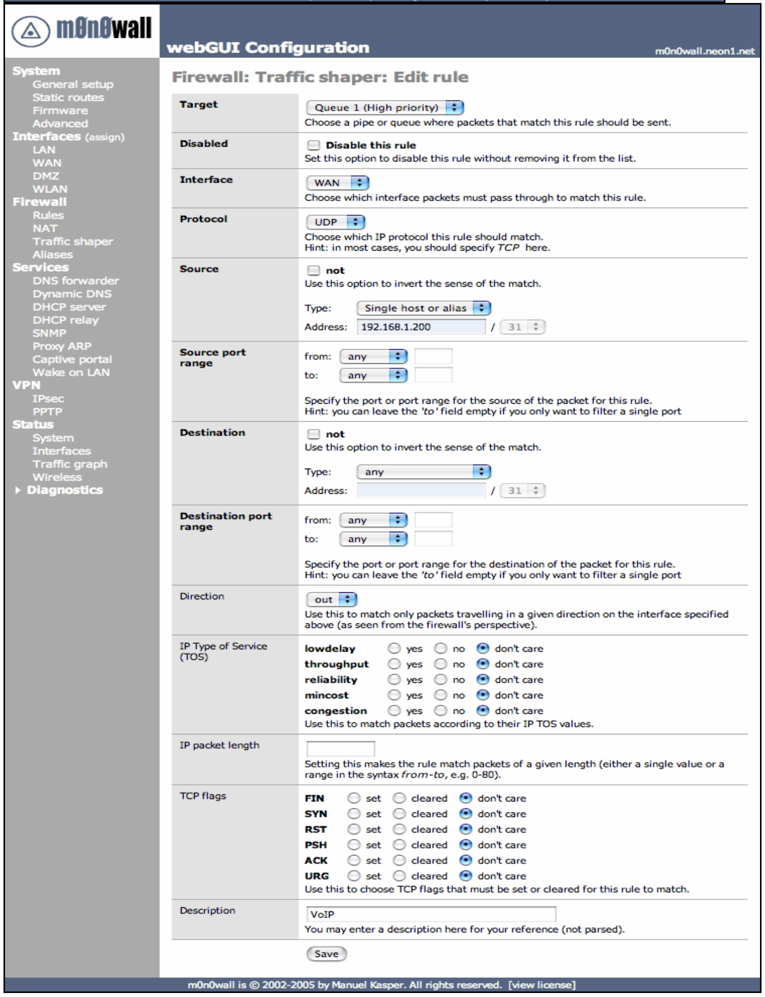

Firewall: Traffic shaper: Edit rule

Target	Queue 1 (High priority) Choose a pipe or queue where packets that match this rule should be sent.
Disabled	□ **Disable this rule** Set this option to disable this rule without removing it from the list.
Interface	WAN Choose which interface packets must pass through to match this rule.
Protocol	UDP Choose which IP protocol this rule should match. Hint: in most cases, you should specify *TCP* here.
Source	□ **not** Use this option to invert the sense of the match. Type: Single host or alias Address: 192.168.1.200 / 31
Source port range	from: any to: any Specify the port or port range for the source of the packet for this rule. Hint: you can leave the *'to'* field empty if you only want to filter a single port
Destination	□ **not** Use this option to invert the sense of the match. Type: any Address: / 31
Destination port range	from: any to: any Specify the port or port range for the destination of the packet for this rule. Hint: you can leave the *'to'* field empty if you only want to filter a single port
Direction	out Use this to match only packets travelling in a given direction on the interface specified above (as seen from the firewall's perspective).
IP Type of Service (TOS)	**lowdelay** ○ yes ○ no ● don't care **throughput** ○ yes ○ no ● don't care **reliability** ○ yes ○ no ● don't care **mincost** ○ yes ○ no ● don't care **congestion** ○ yes ○ no ● don't care Use this to match packets according to their IP TOS values.
IP packet length	Setting this makes the rule match packets of a given length (either a single value or a range in the syntax *from-to*, e.g. 0-80).
TCP flags	**FIN** ○ set ○ cleared ● don't care **SYN** ○ set ○ cleared ● don't care **RST** ○ set ○ cleared ● don't care **PSH** ○ set ○ cleared ● don't care **ACK** ○ set ○ cleared ● don't care **URG** ○ set ○ cleared ● don't care Use this to choose TCP flags that must be set or cleared for this rule to match.
Description	VoIP You may enter a description here for your reference (not parsed).

m0n0wall is © 2002-2005 by Manuel Kasper. All rights reserved. [view license]

References

[1] Telus. Advertisement. June 1 2006. <http://businesscontent.telus.com/webcontent/content/Products/internetData/internetAccess/highSpeed/highSpeedPickPlan.jspe>

[2] Rogers Telecom. Homepage. July 15 2006. <http://www2.rogerstelecom.ca/business/>

[3] Phanse, Kaustubh. Analysis of TCP Performance overAsymmetric Wireless Links. April 25, 2000. <http://fiddle.visc.vt.edu/courses/ecpe6504wireless/projects_spring2000/report_phanse.pdf>

[4] G.711 Protocol Overview. August 1 2006. www.freesoft.org/CIE/Topics/127.htm

[5] Shaw Cablesystems. Advertisement. July 15 2006. <http://www.shaw.ca/en-ca>

[6] Rogers. Advertisement. August 1 2006. <http://www.shoprogers.com/store/cable/InternetContent/internet.asp?shopperID=PP8P30A6N0U59LGLJACSUUK1603Q7CD0>

[7] Cogeco Cable. Advertisement. August 1 2006. <http://www.cogeco.ca/en/residential_shop_for_voip_o.html>

[8] Webopedia. ILEC. August 1 2006. < isp.webopedia.com/TERM/I/ILEC.html>

[9] Aware. ADSL2 AND ADSL2+ THE NEW ADSL STANDARDS. . August 1 2006. <www.dslprime.com/a/adsl21.pdf>

[10] Big Band Networks. DOCSIS 2.0. August 1 2006. <www.bigbandnet.com/documents/BigBand_Networks_Docsis2.pdf>

[11] Webopedia. Fixed Wireless. August 1 2006. <http://www.webopedia.com/TERM/F/fixed_wireless.html>

[12] Answers. Satellite Communications. August 1 2006. <http://www.answers.com/topic/satcom-satellite>

[13] Cisco Systems. Homepage. August 1 2006. < http://www.cisco.com/>

[14] Hewlett Packard. Homepage. August 1 2006. <http://www.hp.com/>

[15] 3COM. Homepage. August 1 2006. <http://www.3com.com/index2.html>

[16] Nortel. Homepage. August 1 2006 < http://www.nortel.com/ >

[17] Webopedia. Latency. August 1 2006. <http://isp.webopedia.com/TERM/L/latency.html>

[18] Cisco Systems. Understanding Delay in Packet Voice Networks. Aug 1 2006. <http://www.cisco.com/warp/public/788/voip/delay-details.html>

[19] Whatis. Jitter. August 1 2006. <http://searchvoip.techtarget.com/sDefinition/0,,sid66_gci213534,00.html>

[20] Wiki. Quality of Service. Aug 1 2006. <http://www.voip-info.org/wiki/view/QoS>

[21] Webopedia. VoIP Packet Loss. August 1 2006.
<http://www.webopedia.com/TERM/V/VoIP_packet_loss.html>

[22] QoVia. How Delay and Packet Loss impact Voice Quality. August 1 2006.
<http://www.qovia.com/resources/PDFs/white%20papers/How%20Delay%20and%20Packet%20Loss%20Impact%20Voice%20Quality%20in%20VoIP.pdf>

[23] Cisco Systems. FIFO Queuing. August 1 2006.
<http://www.cisco.com/en/US/tech/tk543/tk544/tk228/tsd_technology_support_sub-protocol_home.html>

[24] Weighted Fair Queuing. August 1 2006.
<http://www.sics.se/~ianm/WFQ/wfq_descrip/node21.html>

[25] Cisco Systems. Weighted Fair Queueing. August 1 2006.
<http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/120newft/120t/120t5/cbwfq.htm#wp17641>

[26] OpenBSD. Packet Queuing and Prioritization. August 1 2006.
<http://www.openbsd.org/faq/pf/queueing.html>

[27] Cisco Systems. Low Latency Queuing. August 1 2006.
<http://www.cisco.com/warp/public/121/latencyqueueing.html>

[28] IANA. Port Numbers. August 1 2006 <www.iana.org/assignments/port-numbers>

[29] G.711. August 1 2006. <http://www.voip-info.org/wiki-ITU+G.711>

[30] RFC. The Secure Shell (SSH) Protocol Architecture. August 1 2006. <http://www.rfc-archive.org/getrfc.php?rfc=4251>

[31]What Is. Denial of Service. August 1 2006.
<http://whatis.techtarget.com/definition/0,289893,sid9_gci213591,00.html>

[32] RFC. Type of Service in the Internet Protocol Suite. August 1 2006.
<http://rfc.sunsite.dk/rfc/rfc1349.html>

[33] Rhys Haden. IP Datagram. *August 1 2006.* <http://www.rhyshaden.com/ipdgram.htm>

[34] RFC. Definition of the Differentiated Services Field (DS Field)in the IPv4 and IPv6 Headers . August 1 2006. <http://www.rfc.net/rfc2474.html>

[35] Benzedrine. Prioritizing TCP ACKs. August 1 2006. <http://www.benzedrine.cx/ackpri.html>

[36] Phanse, Kaustubh. Analysis of TCP Performance over Asymmetric Wireless Links. August 1 2006. <http://fiddle.visc.vt.edu/courses/ecpe6504-wireless/projects_spring2000/report_phanse.pdf>

[37] RFC. Assigned Numbers. August 1 2006. <http://rfc.sunsite.dk/rfc/rfc739.html>

[38] Ellacoya Networks. Homepage. August 1 2006. <http://www.ellacoya.com>

[39] Packeteer. Homepage. August 1 2006. <http://www.packeteer.com>

[40] Allot. Netenforcer. August 2 2006. <http://www.allot.com/html/products_netenforcer.shtm>

[41] Sandvine. Policy Traffic Switch. August 1 2006. <http://www.sandvine.com/products/policy_traffic_switch.asp>

[42] Packeteer. Packetshaper. August 1 2006. < http://www.packeteer.com/resources/prod-sol/spec_sheet.pdf>

[43] Cisco Systems. NBAR. August 1 2006. <http://www.cisco.com/en/US/products/ps6616/products_ios_protocol_group_home.html>

[44] Cisco Systems. Netflow. August 1 2006. <http://www.cisco.com/en/US/products/ps6601/products_ios_protocol_group_home.html>

[45] Monowall. Homepage. August 1 2006. <http://m0n0.ch/wall/>

[46] Rizzo, Luigi. Dummynet. August 1 2006. <http://info.iet.unipi.it/~luigi/ip_dummynet/>

[47] Mastershaper. Homepage. August 1 2006. <http://www.mastershaper.org/tiki-index.php>

[48] IPTables. Homespage. August 1 2006. <http://www.netfilter.org/projects/iptables/index.html>

[49] Open BSD. Homepage. August 1 2006. <http://www.openbsd.org/>

[50] Bandwidth Controller. Homepage. August 1 2006. <http://bandwidthcontroller.com/>

[51] Softperfect Bandwidth Shaper. Homepage. August 1 2006. <http://www.softperfect.com/products/bandwidth/bandwidth-shaping.htm>

[52] DU Super Controller. Homepage. August 1 2006. <http://www.homeqos.com/>

[53] Inetshaper. Homepage. August 1 2006. <http://inetshaper.com>

[54] Shaw. Advertisement. August 1 2006. <http://www.shawbigpipe.com/html/internet_gateway_big.jsp>

[55] IETF. Internet Protocol Specification. August 1 2006. <www.ietf.org/rfc/rfc0791.txt>

[56] Nelleman, Adam. M0n0wallDocumentationTestBed : TrafficShaperHowTO. August 1 2006. <http://wiki.m0n0.ch/wikka.php?wakka=TrafficShaperHowTO>

[57] IPERF. Homepage. August 1 2006. <http://dast.nlanr.net/Projects/Iperf/>

[58] FPING. Homepage. August 1 2006. <http://www.kwakkelflap.com/fping.html>

[59] NetPERF. Homepage. August 1 2006. <http://www.netperf.org/netperf/NetperfPage.html>

[60] Ethereal. Homepage. August 1 2006. <http://www.ethereal.com/>

[61] Commview. Homepage. August 1 2006 <www.tamos.com/products/commview/>

[62] Comtraffic. Homepage. August 1 2006. < http://www.tamos.com/products/commtraffic/>

[63] Internet Protocol, Version 6 (IPv6). August 1 2006. <www.faqs.org/**rfc**s/**rfc**2460.html>

[64] Monowall. Screenshots. August 1 2006. <http://m0n0.ch/wall/screenshots.php>