NETEQUALIZER

WHITE PAPER

Equalizing Compared to Application Shaping

(Layer-7 "Deep Packet Inspection" Products)

In the following sections, we will discuss equalizing, NetEqualizer's behavior-based shaping algorithms, including what equalizing is and specifically when and where it should be used. We will highlight how it fits into the overall application shaping landscape, and how in many cases equalizing is a great alternative. We will then cover when and where application shaping is used, how it can be used to your advantage, and also when it may not be a good option for what you are trying to accomplish.

What is Equalizing?

Overall, equalizing is a simple concept to understand. *Equalizing is the art form of looking at the usage patterns or traffic "behaviors" on the network, and then when things get congested, robbing from the rich (bandwidth hogs) to give to the poor.*

This *behavior-based approach* to traffic shaping usually mirrors what you would end up doing if you could see and identify all of the traffic on your network, but does not require the labor and cost of classifying everything. Applications such as VoIP, web-based business applications (SaaS, cloud-based apps), Internet browsing, and instant messaging (IM) all naturally receive higher priority, while large downloads and video receive lower priority. As behavior-based shaping does not care what the traffic type (aka "application") is, it also does not need to be updated constantly as applications change.

Once equalizing is in place, it automatically shapes your network when it is congested, using algorithms to implement "instantaneous fairness". The concept of "instantaneous fairness" enables your network to continue providing quick response times to the majority of your users during peak periods of congestion, while restricting bandwidth hogs from consuming your entire network. Low bandwidth users do not have to share the pain of a slow, congested network clogged by larger, network-hogging applications.

Rather than writing hundreds of rules to specify allocations to specific traffic as in application shaping, you can simply assume that when your network is congested, large bandwidth uses need to be curtailed, short quick traffic can run uninterrupted, and be done with it.

Equalizing will keep your network from coming to a standstill and will also increase the load that your network can sustain at peak. During peak periods of congestion on your network, equalizing will apply its proprietary algorithms to shape traffic. Equalizing traffic shaping ensures that a large bandwidth hog does not bring the network to a crawl and that business can continue as usual.

Equalizing Summary

- Suited to Internet or WAN links
- Simple turn-key solution, configure once and done
- All traffic runs unencumbered when network is not congested
- During peak congestion, shapes using fairness-based algorithms
- The most effective for shared Internet trunks
- Real-time reporting by behavior, including bandwidth use, penalties, and P2P traffic
- Real-time protocol and bandwidth graphs
- Behavior-based control of encrypted & unencrypted P2P
- Affordable, low entry cost
- Very little recurring cost or labor
- Supports Net Neutrality

How does Equalizing work?

Let's review equalizing in action. When your network is at peak congestion, equalizing evaluates each inbound and outbound data flow (IP pair) on your network. Data flows vary widely in bandwidth consumed, and will vary from short dynamic bursts, such as when searching a small website, to large persistent flows, as when performing a file download.

When shaping is occurring on your network, equalizing looks at the following factors to intelligently make shaping decisions:
1) How persistent is this data flow?
2) How many active data flows are there?
3) How long has this data flow been active?
4) How congested is the overall network trunk?
5) How much bandwidth is this data flow using, relative to the network trunk size?

These questions are being asked and assessed in real-time, on all inbound and outbound data flows on your network. *As congestion can occur in either direction, we equalize bi-directionally.* Equalizing will make adjustments to your traffic flow, by adding latency to large persistent flows, so that short/bursty data flows receive sufficient bandwidth. This eases congestion in two ways: first, and most obvious, the large flow will reduce its bandwidth consumption, thereby enabling other traffic to get a share of the network pipe. Second, and less obviously, the sender will in many cases "back off" on the data flow rate, which also frees up network bandwidth.

Another subtlety of equalizing is that traffic flows unimpeded when the network has available bandwidth.

> *"Equalizing is the art form of looking at the usage patterns or traffic "behaviors" on the network, and then when things get congested, robbing from the rich (bandwidth hogs) to give to the poor."*
>
> Art Reisman, CTO, APconnections

The implication here is that shaping is not applied when it is not needed. This is an important distinction to remember, as you will see later in our discussion that this is not true of application shaping.

A final distinction to be aware of is that equalizing is not user-based. As we have stated, equalizing logic is applied to individual data flows (IP pairs). This enables it to discriminate across traffic for an individual user. For example, as shown below, a network user sending an email, browsing the web, while on a VoIP call, and downloading a large file will be treated as four (4) separate data flows. During periods of network congestion, the VoIP, email, and browsing sessions will continue seamlessly, while the large file download will be slowed in order to prevent the VoIP call from breaking up, the web session from freezing up, or the email to hang.

Relieving Hidden Node Congestion

Equalizing is the only traffic shaping technology that can relieve hidden node congestion in wireless networks. An independent study by the University of Limerick validated that equalizing technology solves the hidden node issue.

Net Neutrality

And finally, equalizing supports Net Neutrality, and maintains user privacy, as it does not inspect Internet packets.

How accurate is Equalizing?

As equalizing is not inspecting traffic, nor trying to determine the type of application for that traffic, *classification accuracy is not an issue*. The very nature of behavior-based shaping enables traffic to be managed and controlled without worrying about classification.

Equalizing is looking at the size of the traffic and how the traffic behaves on your network. These are the things that really matter when you think about it. If your goal is to optimize your network resources, understanding and responding to the nature of traffic on your network just makes sense. The additional overhead of attempting traffic classification is not required to speed up your network.

In addition, equalizing does a good job of not penalizing traffic unnecessarily. It looks at average traffic behavior over a defined period of time to make sure that traffic is not penalized for short bursts of large bandwidth use.

When should you choose Equalizing?

Equalizing is the right solution when you are trying to shape traffic on a shared Internet trunk, or if you want a simple turnkey way to optimize VoIP and Business Applications on a WAN link.

Internet traffic is highly variable. The application types that will be popular 1-2 years from now may not even exist today. Equalizing is a great choice for shaping Internet trunks, as it does not need to know in advance what types of applications are popular.

Equalizing gets you out of the classification game. *As equalizing does not rely on classification to shape traffic,* you do not have to worry about common application shaping challenges, such as misclassification of traffic, unclassifiable traffic, or only 90% of your traffic getting classified correctly.

Equalizing maximizes the use of your bandwidth, which can help you to defer a bandwidth upgrade and also ensures that you are optimizing your existing bandwidth. Your bandwidth is available for all application types to use as needed. You do not need to allocate portions of your network to each application type, trying to determine in advance what to reserve for each type, which can leave valuable bandwidth unused. Equalizing does use a small amount of overhead in which to shape your network. However, the trade-off is that the user experience at peak is significantly improved.

Equalizing enables you to minimize management of your network pipe. You do not need to spend time keeping up with policy files that capture the latest types of applications. Equalizing works on all types of traffic, regardless of the application. This saves your network administrator time.

Equalizing also saves you money. You no longer need to subscribe to expensive policy file updates. As we believe that bandwidth shaping should be affordable, we have priced our solution fairly, which we believe makes us *the most cost-effective solution on the market.*

Equalizing has also been very successfully implemented on corporate WANs, particularly in situations where a business is looking to provide QoS without having to carve up the network to support defined protocols.

How is P2P Traffic Handled?

The other main cause of Internet gridlock, as well as bringing down routers and access points, is peer-to-peer (P2P) traffic. P2P traffic can be recognized by its propensity to open hundreds or perhaps thousands of small connections to different sources on the Internet.

P2P traffic can be a major source of congestion, as well as be used for illegal activities, such as downloading licensed movie or music content. Therefore, on some networks it is advisable to stop P2P traffic from either the entire network or portions of it.

Equalizing takes care of instances of congestion caused by single-source bandwidth hogs, but does not address P2P directly. Over the years, NetEqualizer engineers have developed additional algorithms to spot connection abuse and avert its side effects.

Consistent with our behavior-based shaping, we look at the behavior of connections to determine if P2P traffic is active on your network. When an IP address, subnet, or the entire network has traffic that is trying to open hundreds or thousands of connections, connection limits will stop opening connections after a predefined limit. This effectively “stops” P2P traffic.

Using Connection Limits, the NetEqualizer is able to deter both encrypted and unencrypted P2P traffic, without the additional overhead of classifying this traffic.

What about Video Traffic?

Today video is a becoming a larger component of Internet traffic. The use of YouTube, Netflix, Hulu, and other video sites to play live streaming video is prevalent. Corporate and commercial hosted video sites for training, teaching, and video conferencing use are also on the rise.

Equalizing supports video use. During periods where your network is flowing freely, equalizing will let video run unencumbered. When your network is congested, video can quickly cripple your network. During peak, equalizing will kick in to ensure that video does not consume all of your bandwidth.

In addition, we offer several options to handle video traffic that can enhance video QoS during peak periods:

Prioritizing Live Streaming Video

Where business critical video is being used, we find that in many cases this is run from a hosted or known server. The NetEqualizer has a low-level routine that easily allows you to give overriding priority to a specific server, whether internal or external to your network. Priority can be configured either temporarily for one-time events, or permanently.

This enables you to ensure that live-streaming video is not equalized during periods of network congestion. Prioritization should be used sparingly, and only for business critical video needs.

The NetEqualizer Caching Option (NCO)

NetEqualizer Caching Option stores a local copy of all port 80 traffic file sizes from 2MB to 40MB. We find that caching primarily targets YouTube videos. However, any type of static content that is frequently accessed will benefit from caching. NCO is integrated with Equalizing, providing a comprehensive bandwidth management strategy. Traffic can be accessed from cache or accessed from the Internet and equalized, as needed.

What about NetEqualizer Reporting?

For equalizing, the key reports are all about viewing real-time traffic on the network. The focus is to help with reviewing whom the bandwidth hogs are, where penalties have been applied, and what the top active data streams are on the network. Traffic exhibiting P2P patterns can also be identified.

Both graphical and tabular reports are available, and display data for all data streams (IP pairs) or selected IPs. Reports display both upload and download bandwidth usage, as well as (coming soon) penalties over time.

As equalizing automatically takes care of providing QoS on the network, reports are informational only, and are not needed to tweak and tune shaping.

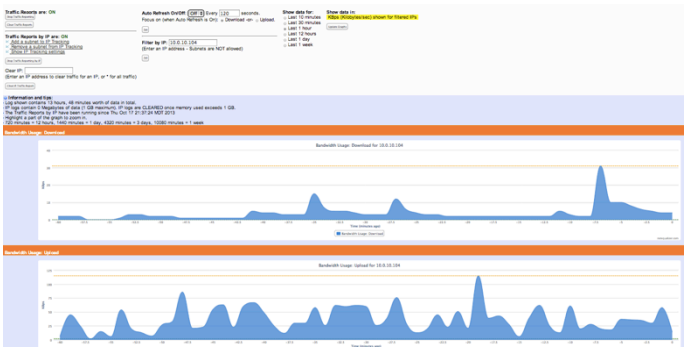

We also offer protocol reports, to help in capacity planning for networks. While protocols are not used for shaping, it can be useful to know how much of each type of traffic is running on the network.

However, in our opinion, while reporting is essential, complicated reporting tools tend to be overkill. Detailed bandwidth monitoring technology is not only more expensive from the start, but an administrator is also likely to spend more time making adjustments and looking for optimal performance. The result is a continuous cycle of unnecessarily spent manpower and money.

What is Application Shaping?

Application Shaping is defined as the ability to identify traffic on your network by type, and then set customized policies to control the flow rates for each particular type. For example, Citrix, AIM, YouTube, and BearShare are all applications that can be uniquely identified. As you are likely aware, all traffic on the Internet travels around in what is called an IP packet. An IP packet can very simply be thought of as a string of characters moving from computer A to computer B. The string of characters

is called the “payload,” much like the freight inside a railroad car. On the outside of this payload is the address where it is being sent. On the inside is the data/payload that is being transmitted. These two elements, the address and the payload, comprise the complete IP packet. In the case of different applications on the Internet, we would expect to see different kinds of payloads.

“At the heart of all current application shaping products is special software that examines the content of Internet packets, performing "deep packet inspection"…

At the heart of all current application shaping products is special software that examines the content of Internet packets, performing "deep packet inspection", as they pass through the packet shaper. Through various pattern-matching techniques, the packet shaper determines in real-time what type of application a particular flow is. It then proceeds to take action to possibly restrict or allow the data, based on a pre-set rule designed by the system administrator.

How Accurate is Application Shaping?

As application shaping needs to examine the content of Internet packets, the question of accuracy comes into play. The challenges with classifying Internet packets are numerous. We will discuss the key challenges in detail below.

Misclassification of Traffic

Traffic can easily be misclassified. For example, the popular peer-to-peer (P2P) application Kazaa actually has the ASCII characters “Kazaa” appear in the payload; and hence a packet shaper can use this keyword to identify a Kazaa application. Seems simple enough, but suppose that somebody was downloading a Word document discussing the virtues of peer-to-peer and the title had the character string “Kazaa” in it. Well, it is very likely that this download would be identified as Kazaa and hence misclassified. After all, downloading a Word document from a web server is not the same thing as the file-sharing application Kazaa.

Unclassifiable Traffic

The other issue that constantly brings the accuracy of application shaping under fire is that some application writers find it in their best interest not be classified. In a mini arms race that plays out everyday across the world, some application developers are constantly changing their signature, and some have gone as far as to encrypt their data entirely.

Yes, it is possible for the makers of application shapers to counter each move, and that is exactly what the top companies do; but it can take a heroic effort to keep pace. The constant engineering and upgrading required has an escalating cost factor. In the case of encrypted applications, the amount of CPU power required for decryption is quite intensive and impractical. We believe that other methods will be needed to identify encrypted P2P.

Application Shaping Summary

- Suited to WAN links
- Must set-up and maintain defined policies
- Pre-defined policies applied at all times
- Not the best fit for shared Internet trunks
  - *Unpredictable traffic types make policies inaccurate*
  - *40% of traffic unclassifiable due to SSL encryption*
- Detailed reporting by application type
- Does not handle encrypted P2P. Expect false positives.
- Expensive. High initial cost
- High ongoing costs for licensing and labor
- Violates Net Neutrality

This is not to say that application shaping doesn't work in some cases or provide some value. So, let's break down where it has potential, and where it may bring false promises. First off, the realities of what really happens when you deploy and depend on this technology need to be discussed.

The Sixty Percent Rule

As of early 2003, we had a top engineer and executive join APconnections direct from a company that offered application shaping as one of their many value-added technologies. He had first-hand knowledge from working with hundreds of customers who were big supporters of application shaping. *The application shaper his company offered could identify 90 percent of the spectrum of applications, which means they left 10 percent as unclassified. So, right off the bat, 10 percent of the traffic is unknown by the traffic shaper.* Since that time, the growth in encrypted traffic has reduced identification down to 60%, as SSL encrypted traffic cannot be identified.

Is this traffic important? Is it garbage that you can ignore? Well, there is no way to know without any intelligence, so you are forced to let it go by without any restriction. Or, you could put one general rule over all of the traffic – perhaps limiting it to 1 megabit per second maximum, for example. Essentially, if your intention was 100-percent understanding and control of your network traffic, right out of the gate you must compromise this standard.

Unclassifiable traffic is a growing problem, and really points out how futile it can be to try to classify traffic. But how accurate can a packet shaper be with the traffic it does claim to classify? Does it make mistakes?

False Positives

There really isn't any reliable data on how often an application shaper will misidentify an application. To our knowledge, there is no independent consumer reporting company that has ever created a lab capable of generating several thousand different application types with a mix of random traffic, and then took this mix and identified how often traffic was misclassified. Yes, there are trivial tests done one application at a time, but misclassification becomes more likely with real- world complexity and diverse application mixes.

From our own testing with application technology freely available on the Internet, *we discovered false positives could occur up to 25 percent of the time.*

Obviously, commercial packet shapers do not rely on free technology like open source, and they may improve upon it.

So, if we had to estimate based on our experience, perhaps 5 percent of Internet traffic will likely get misclassified. *This brings the overall accuracy of packet shaping down to 55 percent* (combining the traffic they don't claim to classify with an estimated error rate for the traffic they do classify).

Constantly Evolving Traffic

Our sources say that *70 percent of customers that purchased application shaping equipment were using the equipment primarily as a reporting tool after one (1) year.* This means that they had stopped keeping up with shaping policies altogether, and were just looking at the reports to understand their network. After one year, they were doing nothing proactive to shape their network traffic.

This is an interesting fact. From what we have seen, many people are just unable, or unwilling, to put in the time necessary to continuously update and change their application rules to keep up with evolving traffic.

The reason for the constant changing of rules is that with traditional application shaping you are dealing with a cunning and wise foe. For example, if you notice that there is a large contingent of users using Bittorrent, and you put a rule in to quash that traffic, within perhaps days those users will have moved on to something new: perhaps a new application or encrypted P2P.

If you do not go back and re-analyze and reprogram your rule set, your packet shaper slowly becomes ineffective.

And finally, lest we forget, application shaping is considered by some to be a violation of Net Neutrality, due to the very nature of packet inspection.

When should you choose Application Shaping?

WAN Optimization

There is a large set of businesses that use application shaping quite successfully, along with other technologies. This area is WAN optimization. Thus far, we have discussed the issues with using an application shaper on the wide-open Internet where the types and variations of traffic are unbounded.

In a corporate environment with a finite set and type of traffic flowing between offices, an application shaper can be set up and used for WAN optimization with reliable results. We have also achieved equal and sometimes better results with a NetEqualizer head-to-head in a WAN environment.

What about Application Shaping Reporting?

An industry-leading packet shaper brings visibility to your network and a pie chart showing 300 different kinds of traffic. *Whether or not the tool is practical or accurate over time isn't often brought into the buying decision.*

It is human nature to want to see and control what takes place in your environment. Finding the best tool available to actually show you what is on your network, and the perceived ability to contain it, plays well with just about any CIO or IT Director on the planet. The downside of detailed reporting over a simple heuristic solution is that it can create data confusion. We have addressed this subject in our article, "The True Price of Bandwidth Monitoring".

As the cost of bandwidth continues to fall, the question becomes how much a CIO should spend to analyze a network. This is especially true when you consider that as the Internet expands, the complexity of shaping applications grows. As bandwidth prices drop, the cost of implementing such a product is either flat or increasing. In cases such as this, it often does not make sense to purchase a $25,000 bandwidth shaper to stave off a bandwidth upgrade that might cost an additional $200 a month.

Even if it can only accurately report 55 percent of the actual traffic, isn't this useful data in itself? Yes and no. Obviously analyzing 55 percent of the data on your network might be useful, but if you really look at what is going on, it is hard to feel like you have control or understanding of something that is so dynamic and changing.

By the time you get a handle on what is happening, the system has likely changed. Unless you can take action in real-time, the network usage trends (on a wide-open Internet trunk) will vary from day-to-day.[1] It turns out that the most useful information you can determine regarding your network is an overall usage pattern for each individual. The goof-off employee/user will stick out like a sore thumb when you look at a simple usage report, since the amount of data transferred can be 10-times the average for everybody else. The behavior is the indicator here, but the specific data types and applications will change day-to-day and week-to-week.

1 The exception is a corporate WAN link with relatively static usage patterns.

Conclusion

Hopefully you now have a good understanding of equalizing, and how it is different from application shaping. If you come from the application shaping world, equalizing takes some getting used to at first, and then our customers rave about their faster networks, and love the low maintenance "set it and forget it" aspect of our model.

We hope that this white paper has given you the facts that you need to make an informed decision about what direction is right for you - application shaping or equalizing.

If you have additional questions, please feel free to contact us via email: sales@apconnections.net or call us at 303.997.1300 x103.

Summary Table: Equalizing Compared to Application Shaping

Equalizing	Application Shaping
Suited to Internet or WAN links *Good for links where traffic patterns are variable (Internet), or you prefer not to define (WAN).*	Suited to WAN links *Good for static links where traffic patterns are constant (WAN)*
Simple turn-key solution, configure once and done	Must set-up and maintain defined policies
All traffic runs unencumbered when network is not congested. During peak congestion, shapes using fairness-based algorithms.	Pre-defined policies applied at all times. Policies applied whether network is congested or not.
The most effective for shared Internet trunks	Not the best fit for shared Internet trunks *Unpredictable traffic types make policies inaccurate* *40% of traffic unclassifiable due to SSL encryption*
Real-time reporting by behavior, including bandwidth use, penalties, and P2P traffic Real-time protocol and bandwidth graphs	Detailed reporting by application type
Behavior-based control of encrypted and unencrypted P2P	Does not handle encrypted P2P Expect false positives
Affordable. Low entry cost	Expensive. High initial cost
Little or no recurring cost or labor	High ongoing costs for licensing and labor
Supports Net Neutrality	Violates Net Neutrality

About APconnections, Inc.

APconnections is an innovation-driven technology company that delivers best-in-class network traffic management solutions to give our customers better networks, with zero maintenance, at the best prices. We specialize in turnkey bandwidth shaping and intrusion prevention system (IPS) appliances. APconnections is based in Lafayette, Colorado, USA. We released our first commercial offering in July 2003, and since then thousands of customers all over the world have put our products into service. Today, our flexible and scalable solutions can be found in many types of public and private organizations of all sizes across the globe, including: Fortune 500 companies, major universities, K-12 schools, and Internet Providers on six (6) continents.

