


# AppEx IPEQ (IP End-to-End QoS)

Accurate, Bidirectional, Prioritized and Fair  
Bandwidth Regulation

AppEx Networks Corporation



With the rapid growth of the network services, applications and personal devices, the request for higher bandwidth has outgrown what the network access providers could offer almost exponentially. Only a few years ago, most people were just using the Internet for emails and plain web browsing, and the majority of the homes owned just a single computer with network access. Nowadays, the Internet is heavily loaded with the traffic of High-Definition media, file sharing, IM, VoIP, social media, network gaming and a lot more. Many homes and most businesses have computers, laptops, tablets, netbooks, phones, game consoles, set-top boxes, etc., accessing the network. And good portion of them are accessing the Internet from wireless providers. Many of such apps and devices demand high bandwidth for better user experiences and efficiency. The quest for better flow control / bandwidth management is now stronger than ever before.

## Flow Control

Here the term “Flow Control” is defined as the process of managing the data rate of the inbound and outbound traffic to properly satisfy the service needs of the flows in the traffic.

Under this model, operators of the Flow Control devices configure a set of traffic parameters, such as bandwidth, latency, priority, etc., to each of the flow groups in the traffic. These traffic parameters are typically converted from the Service Level Agreements (SLAs) between the operators and the clients of the services. The Flow Control devices will massage the traffic following the parameters defined.

Overall, the process of Flow Control is typically divided into two phases:

- **Flow Classification:** this is the first phase of the Flow Control, where the traffic flows are classified into different groups called Aggregations. The traffic parameters are applied to the Aggregations.
- **Flow Treatment:** in this second phase, the flows within each Aggregation are processed according to the traffic parameters configured.

Queuing algorithms are often employed for the flow treatment regardless of the classification mechanisms used. Such algorithms are usually very effective for outbound traffic control. For the inbound traffic, however, because the received packets already have occupied the WAN bandwidth, there is little that the queuing algorithms can do to instantly control the data rate. The typical mitigation is to discard the packets, causing the feedbacks of the protocols to slow down the senders themselves. Since the inbound bandwidth has already been occupied, discarding packets means wasting bandwidth. In a typical corporation deployment, or when P2P applications are involved, there will often be thousands or even hundreds of thousands of flows actively running. Under such conditions there have to be large volumes of packet drops to suppress the inbound traffic. Statistics show that to effectively control the inbound flows the overall traffic usually has to suffer as much as 40% bandwidth loss. All packet losses increase the total cost of the clients or the operators.

Categorized by the flow classification methods, we often see Layer-4 or Layer-7 Flow Control devices in the market, based on the ISO OSI 7-Layer Model. Layer-4 Flow Control is a traditional technology that classifies the traffic based on the layer-3 (IP, etc.) and layer-4 (TCP/UDP, etc.) protocol headers, usually

the 5-tuple of the TCP or UDP flows. Such a classification in the real world is usually too primitive and static to capture the characteristics of the network applications, since most majority of the application allows the ports to be redefined. Even if the application itself listens on a fixed port number, NAT can easily translate it into another via port mapping.

Layer-7 Flow Control performs Deep Packet Inspection (DPI) to capture the characteristics of the contents of the flows. Therefore in most circumstance it is able to accurately recognize the network applications of the flows, regardless of whether they are using standard port numbers or not. However, Layer-7 classification is not silver bullet either. Because of the sheer huge number of the applications, and the variety of traffic contents from one app to another, even different versions of the same app, such classifications usually come with big databases of the content signatures. Furthermore, the signature databases need constant updates when new apps or new versions appear, adding to the total cost of ownership of the Layer-7 Flow Control devices. Also, if the traffic is encrypted, such as in SSL, IPSec enabled corporate networks and most VPN tunnels, the contents may be impossible to decrypt and deep inspection simply does not work.

## AppEx IPEQ (IP End-to-End QoS)

AppEx Networks designed and implemented its own advanced Flow Control algorithms named IPEQ, IP End-to-End QoS, to effectively address the hard problems in both flow treatment and classification.

The following diagram shows the logical composition of the IPEQ:

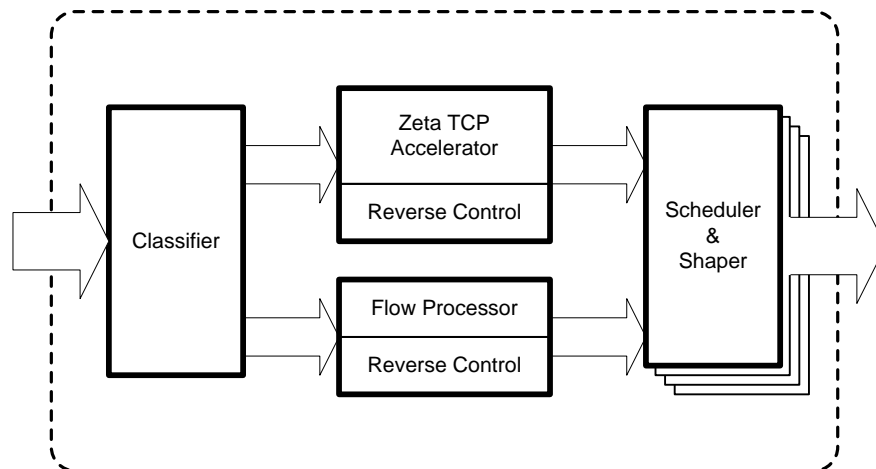


Figure 1: AppEx IPEQ

## Flow Treatment

### Inbound Traffic Shaping

Although capable of shaping the outgoing traffic well, the traditional Flow Control algorithms and implementations are largely, if not completely, ineffective in dealing with the traffic coming in.

Unfortunately inbound bandwidth is what the majority of the users care about most. The typical approach to handle this is to use a Policer or similar algorithm to drop the packets, allowing a certain amount of burst. Such solutions penalize the incoming traffic directly. However, the drawbacks are:


- The bandwidth has already been taken before the packets are discarded. So dropping packets doesn't immediately contribute to the inbound bandwidth control. In our tests with the popular P2P applications, the packet drop could reach above 30% due to incoming policing.
- Discarding the packets may not be immediately fed-back to the sender. It may take the senders a few round trips to start responding to such a signal.
- For protocols with Congestion Control features, for instance, TCP, packet losses may be taken as the indications of Congestion and the senders may lower the sending rate dramatically in response. This results in a large bandwidth waste in over-reacting to the packet losses and retransmitting the data that have been discarded intentionally.
- The traffic policing is very rigid when combined with classifications on priority control, etc. For instance, if a low priority is assigned a certain small portion of the total bandwidth by the policer to guarantee the premium treatment of the higher priority traffic, then even if there are only very low volumes of the higher priorities, the low priority traffic will still not be able to utilize the spare bandwidth.

AppEx IPEQ takes a completely new approach against such a problem, which is called Reverse Control. The essence of Reverse Control algorithms is to avoid packet discard in the first place and control the sender's transmission speed directly from the protocols' feedback mechanism. By maximally reducing the packet losses, AppEx IPEQ is able to protect against bandwidth waste and at the same time achieve smoother conditioned traffic flows, enhancing the end user experiences.

As shown in figure 1, both Zeta-TCP and the flow processing modules have Reverse Control function included. The processing of TCP and other protocol traffic is slightly different. They both work closely with the Scheduler/Shaper to control the inbound flows.

For TCP, we employ an ACK-based algorithm to intelligently and adaptively control the senders' sending rate. For other protocols, especially those encapsulated in UDP, we have studied a large number of network applications and protocols, and come up with an algorithm to intelligently shape the feedback traffic and indirectly achieve the rate control over the inbound flows.

It is worth noting that for TCP there are already a very small number of network devices that perform the similar inbound rate control. Even a few applications managed to achieve a similar result from the socket level. Essentially these approaches are all TCP-Window based, which leverages the adaptive nature of the TCP protocol and imposes a hard throttle on the flows. As a result, the TCP-Window based algorithms all end up with unfavorable spiky/bursty flows, adding more loads to the downstream router queues and causing packet losses. In contrast, AppEx IPEQ algorithm uses the ACK-based approach which is technically more difficult but achieves very nicely shaped flat traffic flows, as shown in the graph of figure 2, which was drawn from the real test data. As the graph indicates, the inbound spikes of the TCP-Window based inbound control struck as high as 4 times the set inbound rate.

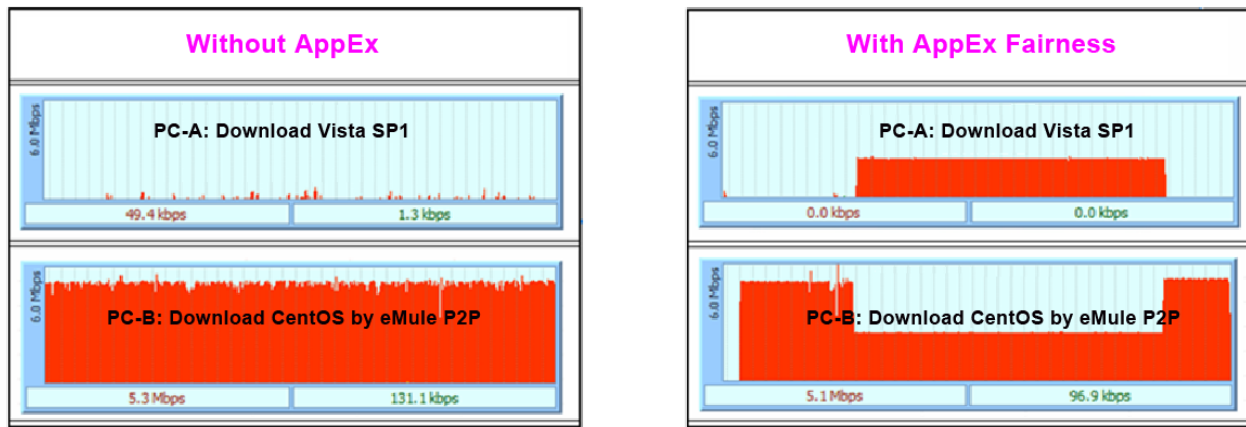


**Figure 2: Inbound Control Comparisons (512 Kbps)**

In figure 2, the curve of the TCP flow with the Policer inbound bandwidth limit is also shown as the blue line. Due to the Policer drop the TCP flow slashed its Congestion Window and retransmitted. The resultant bandwidth was merely about 50% of the Policer’s inbound rate (512 Kbps). The curve also includes the retransmitted packet which added to more wasted bandwidth.

### Fairness

Built upon the accurate inbound and outbound shaping, AppEx IPEQ is able to guarantee bandwidth fairness among Aggregations *of the same priority*. As defined in the previous section, the Aggregations are the result of the classification. We will also explain later in this whitepaper that some Aggregations can also be defined differently under different deployment scenarios. Hence, for instance, AppEx IPEQ is able to achieve fairness among different hosts on the LAN side, or among different applications in the same computer.



**Figure 3: AppEx Host Fairness**


The above figure shows the difference with and without AppEx IPEQ's fairness feature in a real-life deployment test. On the left side, without AppEx IPEQ, the eMule P2P download traffic to PC-B overwhelmed the downlink. When PC-A started downloading Windows Vista SP1 package, it could barely get any bandwidth with its single TCP flow. While on the right side with AppEx host fairness, when PC-A started downloading, PC-B's total bandwidth was suppressed to 50% of the WAN speed so that PC-A would get its fair share.

The AppEx inbound shaping plays an important role in the fairness case above. Without the smoothly shaped inbound traffic, the single TCP flow of PC-A will incur unnecessary packet losses due to the spikes in PC-B's P2P's many flows and those of its own. Had that happened, PC-A would have barely reached 30% of the total bandwidth and PC-B would have still been able to occupy more than 70% of it because of the large number of P2P flows.

## Prioritization

AppEx IPEQ uses a priority number to determine which flows should be processed as early as possible and which are ok be delayed. The higher the priority number, the more important the flow is and the faster the data in it will be processed. The priorities of the inbound and outbound directions of a flow can be separately defined.


Flows of the same priority will be processed fairly. Between different priorities we define the Bandwidth Guarantee and Bandwidth Limit. Bandwidth Guarantee allows the traffic of a priority to use at least up to the given percentage even when there are higher priority flows actively running. Bandwidth Limit is the maximum bandwidth allowed for the total traffic of a priority. All these parameters can be individually configured for each priority.



**Figure 4: Bandwidth Guarantee – Priority 1: 20%, Priority 2: 30%**

Figure 4 illustrates the bandwidth guarantee with flows of different priorities. The red curve was a priority 1 flow, with 20% bandwidth guarantee. The green curve was a priority 2 flow, with 30% bandwidth guarantee. And the blue curve was a priority 3 flow. Without any guarantee the high priority flows could've totally taken over the bandwidth of the low priority ones. Since the low priority flows (1 and 2) both had guarantees, when the blue flow started, it was only be able to take about 50% of the total bandwidth, while the red and green flows was able to keep 20% and 30% bandwidth respectively. Because the green flow had higher priority than the red, without the blue flow traffic it was able to occupy about 80% of the bandwidth while still leaving 20% to the red.

Had all the three flows been of the same priority, the bandwidth would have been equally divided among the 3 due to the fairness within the same priority.



**Figure 5: Flows with the Same Priority (Fairness)**

## Flow Classification

AppEx IPEQ currently supports the following classification methods:

- Layer-4 classification.
- Layer-7 classification.
- Behavior-based classification.
- Dynamic and static classifications.

As explained earlier, layer-4 and layer-7 classifications are available in most of the Flow Control devices today. Although very useful for a very large category of applications, they both have their own problems, being either overly rigid (layer-4), or totally dependent on the often-nonstandard implementation details of the applications. In reality, people are gradually realizing that DPI based classifications are

intrusive, expensive and performance-limiting. To make it worse, some P2P applications are now capable of disguising their data streams to be completely indistinguishable from normal web-browsing traffic content-wise, posing greater challenges to the DPI solutions.

AppEx IPEQ is one of the few solutions that provide **Behavior-Based Classification**. Unlike the layer-4/7 classifications that inspect the specific fields of the packets, behavior-based classification looks into the patterns of the traffic as a whole, mapping the characteristics of the packets and the correlations between a series of packets into different traffic types (Aggregations). With behavior-based classification, we can easily identify, for example, a P2P or a VoIP application without having to dig into the details of each individual implementation, such as Bit Torrent, eDonkey, Skype, Live Messenger, etc. Under many circumstances, the behavior-based classification can be more efficient and effective than the layer-7 classification because of its adaptive and intelligent nature. From a larger aspect, since Flow Control itself is to regulate the behavior of the traffic, classifying by the traffic pattern fits the bill better than classifying by the layer-4/7 contents. The behavior-based classification is a more viable solution in the long run.

Aggregations are the results of the flow classifications. **Dynamic and Static Classifications** are concepts orthogonal to the layer-4/7 and behavior-based ones above, which dictate how the Aggregations are created. The static classification is straightforward, meaning an Aggregation is created for the group of flows that matches a certain rule or pattern. The dynamic classification, on the other hand, means an Aggregation is created for each group of flows that bears a different value for a given rule variable.

Dynamic classification allows a whole new set of results to be created. For example, dynamic classification on source IP address will create different Aggregations for flows from different source IP addresses. Each Aggregation represents the group of flows with the same source IP address. This is the basis of the Host Fairness feature.

The rich set of classification methods provided by AppEx IPEQ facilitate a wide variety of traffic regulation schemes to be deployed at the minimal operational cost. Combined with the unparallel flow treatments, AppEx IPEQ is able to tackle complicated deployment scenarios, offering bidirectional, accurate end-to-end Flow Control with minimum packet losses and fairness and prioritization among the flows.

## Use Scenarios

AppEx IPEQ is capable of regulating the traffic and improving the end users network experiences in a wide variety of scenarios thanks to its bidirectional shaping, feature-rich Flow Control methods and scalability and versatility in its design.

The major advantage of AppEx IPEQ lies within its inbound shaping. As illustrated earlier, the traditional inbound rate control causes bandwidth loss. This is the very reason that many customers found that some network applications became unbearably slow after the deployment of other Flow Control devices. Very often they were force into buying more bandwidth from the ISPs. AppEx's inbound traffic shaping

will relieve the users' pain without losing bandwidth. This protects the customers' investment and saves bandwidth cost.

## Businesses and Organizations


Today's businesses heavily rely on the network for efficiency in their operations. At the same time most of them also provide Internet accesses to their employees at work. Competition of the limited network resources and the high cost of more bandwidth subscriptions constantly make the IT struggle, especially for cross-site accesses among different branches.

With AppEx IPEQ, the network traffic can be easily classified so that crucial business operational applications, such as ERP/CRM, production database, source control, etc., can enjoy higher priorities with premium treatment. At the same time the rest of the bandwidth is still available to the noncritical applications, fairly distributed.

## Schools and Public Network Providers

In the schools (especially colleges and universities) and other similar environments, the network applications vary greatly due to all sort of different needs. And the IT usually has little control over what applications the end users may use. Very often the very few P2P-based applications occupy most of the bandwidth, leaving the majority of the decent / moderate network applications suffer. The same situations also happen in certain businesses such as Internet Cafes. While such businesses often have the control over what applications can be installed in their computers, for business reasons they usually have to provide P2P applications as well.

The P2P applications use a wide variety of protocols, and in the schools there may be brand new bandwidth monsters out of the creative free minds. In such environments AppEx's behavior-based classification with host fairness is the perfect cure.



**Figure 6: P2P Traffic Control**

The above figure shows the bandwidth usage before and after the deployment of AppExIPEQ for a customer. The bandwidth subscribed from the ISP was 30Mbps. Before using AppEx IPEQ, the P2P applications were able to take 18.6Mbps bandwidth while HTTP was able to occupy 0.37Mbps. Then the

customer deployed IPEQ, identifying P2P traffic using behavior-based classification and limiting the P2P traffic up to 20% of the total bandwidth. After such configurations the HTTP was able to shoot up to 20.2Mbps while the P2P traffic was restricted to 6.3Mbps, roughly 20% of the total bandwidth.

## Home Use

A lot more homes now own more than one computer or networked digital devices than just a few years ago. As a result the bandwidth competition within a home emerges. In the majority of the scenarios the preferred solutions would be to have all the devices share the bandwidth fairly and prevent the P2P applications from taking over all the bandwidth. Also there are a few applications that the home users usually wish to have higher priorities than the others, such as VoIP calls, online video streaming, online gaming, etc.

AppEx IPEQ's scalability allows it to be easily built into economic small home routers with limited hardware resources. The behavior based classification is capable of identifying P2P, VoIP and other traffic flows. With AppEx host fairness and a few other predefined Flow Control rules the router will be able to serve the home use scenarios very well.

## Personal Computers and Smart Devices

It may have seemed unlikely a necessity for the computer OS to perform any kind of advanced Flow Control. However, with the rapid development of high-quality media streaming, network gaming, VoIP and live communications, etc., the demand for desktop Flow Control is looming. It is gradually becoming an issue that bandwidth demanding applications are competing bandwidth with large file downloads. Indeed, Microsoft has already realized this issue and provided a service named "Background Intelligent Transfer Service" (BITS) in recent Windows OS that "facilitates prioritized, throttled, and asynchronous transfer of files between machines using idle network bandwidth". However, without the improvement from the protocol level, BITS is merely a cosmetic deferred transmission service that doesn't fully utilize the available bandwidth. Besides only giving the programmers a choice of using a low-priority service, the applications that use BITS must be coded using BITS API, which is Microsoft proprietary and focused on file transfer, very restrictive and different from the socket APIs. Therefore BITS is only able to benefit a very small number of applications.

With the upcoming IE version 9, Microsoft is adding HTML5 support and building hardware accelerations into IE9 to speed up media rendering. The signal is very clear that Internet browsing will be more bandwidth-hunger than ever before. AppEx IPEQ can dramatically improve the end users' network experiences.

AppEx IPEQ is already built in the AppEx Application Bandwidth Controller for PC installation. In PC, the dynamic host-Aggregations are redefined as process-Aggregations. This redefinition enables the **Application Fairness** feature so that no processes will be able to deprive others of network access by hogging the bandwidth. This fairness feature is configuration free, which means no user interactions necessary.

Other than application fairness, the PC users may also define the priorities for different applications to make them work even better when sharing the bandwidth with other applications. For example, the network games, VoIP and online video applications can be marked high-priority to ensure the best flow services are performed for these applications for the best user experiences. The users may also choose to change the applications' priority on-the-fly to speed up or slow down at any time.

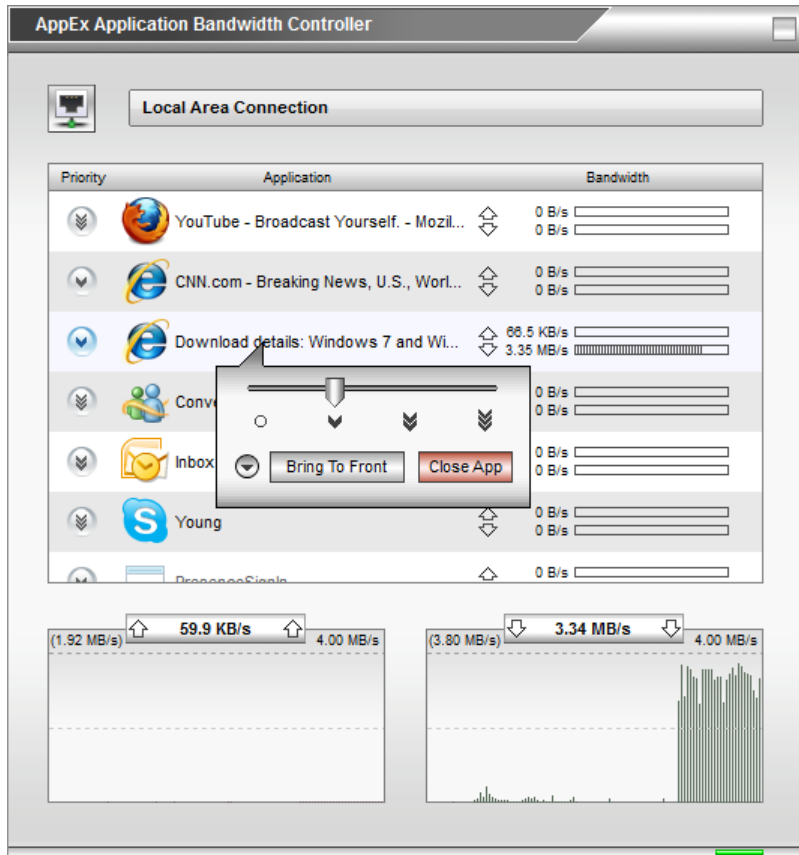


Figure 7: AppEx Application Bandwidth Controller

## Conclusions

The bandwidth competition is the major factor affecting the users' network experiences. Flow Control is the technology to address such issues by properly regulating the traffic through the limited bandwidth to maximize the end-users overall experience.

AppEx IPEQ distinguishes itself from the other Flow Control technologies in its unique inbound shaping which is capable of achieving the similar accurate results as the outbound shaping, while minimizing packet losses and bandwidth waste. It combines the bidirectional accurate shaping, fairness and prioritization with its rich set of classification method, making a very powerful and flexible Flow Control system. Its unparallel scalability and portability allows it to be deployed in a very wide variety of scenarios.