WHITE PAPER

APPEX NETWORKS

AppEx IPEQ™ (IP End-to-End QoS)

Bidirectional, Prioritized and Fair Bandwidth Regulation

Abstract

With the rapid growth of applications and network services being delivered to personal devices, the requirement for higher bandwidth has outgrown what the Internet service providers can make available. Only a few years ago, Internet use was primarily for emails and basic web browsing and the majority of the homes had just a single computer with network access. Today, the Internet is heavily loaded with rich media, file sharing, IM, VoIP, social media, network gaming and more. Most homes and businesses now have multiple computers, laptops, tablets, netbooks, phones, game consoles, set-top boxes, etc.; all accessing the network simultaneously. In addition, a rapidly growing amount of Internet access is coming from mobile networks. All of these applications and devices require high bandwidth for the best user experience. The need for better bandwidth management is stronger now than ever before.

This whitepaper will examine the need for end-to-end quality of service (QOS) in order to meet these bandwidth management challenges.

Flow Control

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 all traffic flows.

Under this model, operators of the flow control devices configure a set of traffic parameters such as bandwidth, latency, priority, etc., for each type of traffic. These traffic parameters are typically obtained from the service level agreements, (SLAs) between the operators and the users of the services. The flow control devices will manage the traffic flow using the defined parameters.

Overall, the process of flow control is typically divided into two phases:

- **Flow Classification** – this is the first phase of flow control where the traffic flows are classified into different groups called Aggregations. The traffic parameters are applied to each Aggregation

- **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. These algorithms are usually very effective for outbound traffic control. For the inbound traffic, however, because the received packets have already occupied the WAN bandwidth, there is little that the queuing algorithms can do to control the inbound data rate. The typical fix is to discard the packets randomly, causing the protocol's normal feedback mechanisms to slow down the senders themselves. Since the inbound bandwidth has already been used, discarding packets means wasting bandwidth. In a typical corporate 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 would 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 loss increases the total cost of the clients and 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. Such a classification in the real world is usually too primitive and static to capture the characteristics of the network applications, since the majority of the applications allow 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 flow content. Therefore under most circumstances it is able to accurately recognize the application flows, regardless of whether they are using standard port numbers or not. However, Layer-7 classification is not a silver bullet either. Because of the sheer number of applications, and the variety of traffic content from one app to another (even different versions of the same app), such classifications usually come with large databases, which contain content signatures. 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 with SSL, IPSec, or other VPN tunnels, the content may be impossible to decrypt and deep inspection simply does not work.

Figure 1: Logical View of IPEQ™

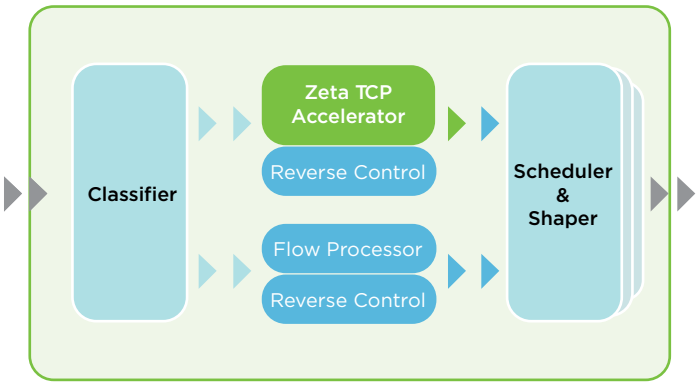

AppEx IPEQ (IP End-to-End QoS)

AppEx Networks designed and implemented its own advanced Flow Control algorithms to effectively address the hard problems in both flow treatment and classification. These algorithms are part of 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 incoming traffic. Unfortunately inbound bandwidth is what the majority of the users care about the most. This would include file downloads, video streaming, VOIP, and application interaction (email, etc.) to name a few. The typical approach to handle this is to use a bandwidth "Policer" or similar algorithm to drop the packets, allowing a certain amount of burst. Such solutions penalize the incoming traffic directly. Following are some of the drawbacks of this approach:

- The bandwidth has already been used before the packets are discarded so dropping packets doesn't immediately contribute to the inbound bandwidth control. In our tests with some popular P2P applications, the packet drops could go as high as > 30% due to incoming policing.
- Discarding the packets may not be immediately fed-back to the sender. It may take the sender a few round trips to start responding to the feedback and slow down.
- For protocols with Congestion Control features like TCP, this forced packet loss may be taken as an indication of Congestion and the senders usually overreact by lowering the sending rate dramatically in response. This overreaction to packet loss wastes bandwidth and causes more re-transmissions because the packets have been intentionally discarded.
- Traffic policing is very rigid when combined with classifications on priority control, etc. For instance, if a low priority flow is assigned a certain small portion of the total bandwidth by the Policer to guarantee the premium treatment of the higher priority traffic flows, 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 in solving such a problem, called Reverse Control. Reverse Control avoids discarding packets in the first place and instead controls the sender's transmission speed directly from the protocols' feedback mechanism. By reducing the packet loss, AppEx IPEQ is able to protect against wasting bandwidth and at the same time smooths traffic flows, which improves the end user experience.

As shown in figure 1, both Zeta-TCP and the flow processing modules have Reverse Control 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, AppEx employs an ACK-based algorithm to intelligently and adaptively control the sending rate of the inbound traffic. For other protocols, especially those encapsulated in UDP, AppEx has studied a large number of network applications and protocols, and has come up with an algorithm to intelligently shape the feedback traffic and indirectly achieve the rate control of the inbound flows.

There have been other attempts by some network device manufacturers to accomplish this inbound rate control with mixed results. These methods use standard TCP feedback mechanisms with hard throttles, which can cause spikes and packet loss. Essentially these approaches are all TCP-Window based, which leverage the adaptive nature of the TCP protocol imposing a hard throttle on the flows. As a result, the TCP-Window based algorithms all end up with spiky/bursty flows, adding more load to the downstream router queues and causing packet loss. In contrast, the AppEx IPEQ algorithm uses an ACK-based approach, which is technically more difficult but achieves nicely shaped flat traffic flows, as shown on the graph in figure 2. This graph was drawn from real test data. The graph shows the inbound spikes caused by the "hard" throttles on inbound traffic.

Figure 2: Inbound Control Comparisons (512 Kbps)

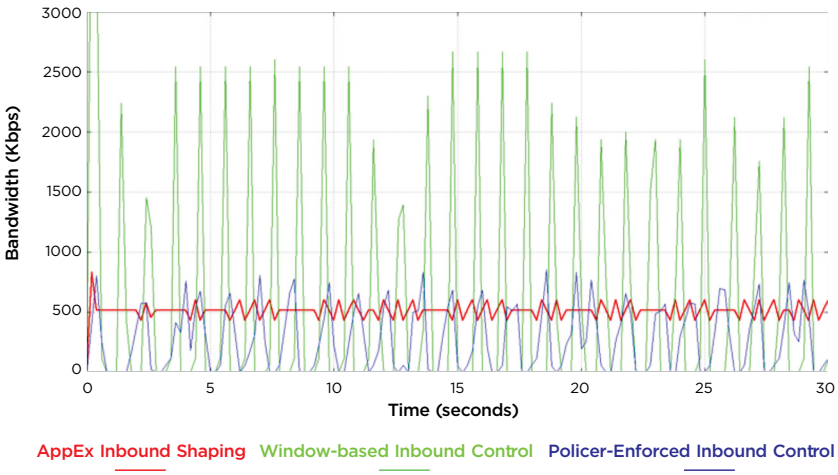

In figure 2, the blue line in the graph above depicts the curve of the TCP flow with the hard throttle (Policer) inbound bandwidth limit. Because the Policer drops packets to control the inbound rate, the TCP flow overreacts, slashing its Congestion Window and retransmitted the dropped packets, which caused spikes and wasted bandwidth. This also results in much lower throughput (less than 50% utilization). The AppEx inbound shaping (red line) is much smoother and more consistent.

Fairness

Because of 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 classifying the flows. We will also explain later in this whitepaper that some Aggregations can also be defined in different ways under different deployment scenarios. For this reason 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

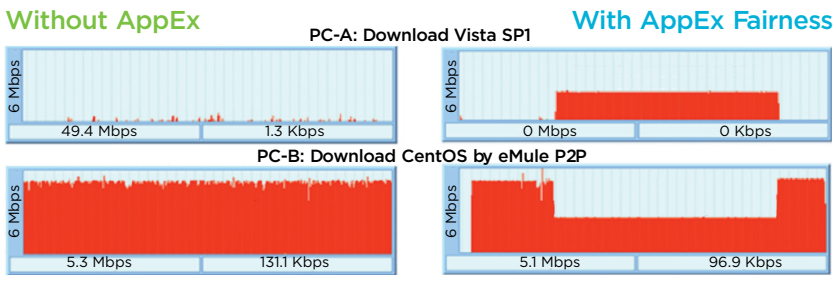

Figure 3 above shows the difference with and without AppEx IPEQ's fairness feature in a real-life deployment test. On the left side, without AppEx IPEQ, eMule P2P downloads traffic to PC-B overwhelming the downlink. When PC-A started downloading Windows Vista SP1 package, it could barely get any bandwidth with its single TCP flow. While on the right side with AppEx host fairness, when PC-A (higher priority) started downloading, PC-B's total bandwidth was suppressed to 50% of the bandwidth so that PC-A would get its fair share.

AppEx inbound shaping plays an important role in the fairness case above. Without IPEQ shaping the inbound traffic, the single TCP flow of PC-A would be starved for bandwidth by

PC-B's P2P's multi-flow dominance. In this scenario PC-A would have barely reached 30% of the total bandwidth and PC-B would have still been able to occupy more than 70% 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.

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

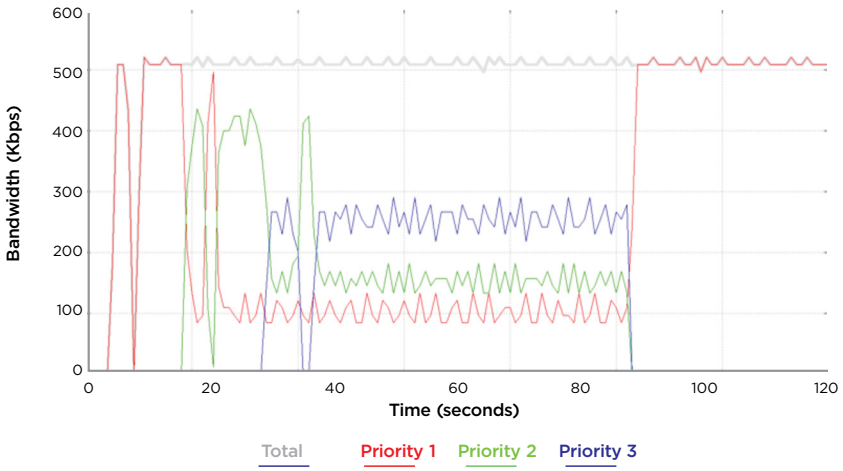

Figure 5: Flows with the Same Priority (Fairness)

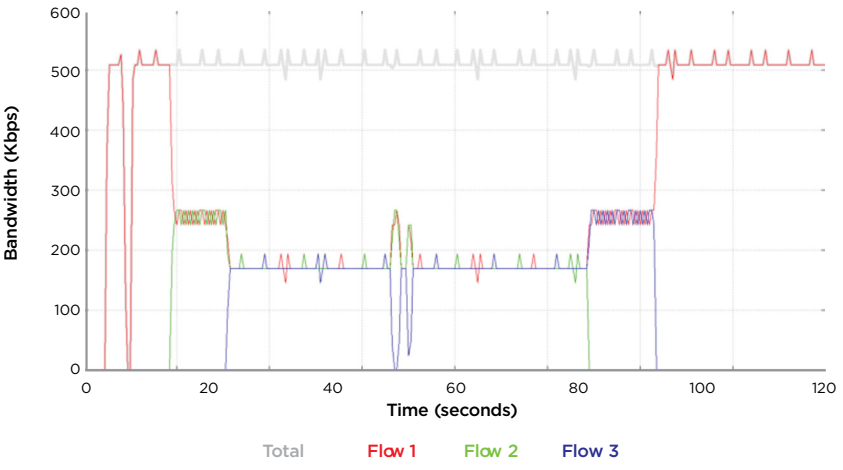

Flows of the same priority will be processed fairly within their own group. For each priority number we define the Bandwidth Guarantee and Bandwidth Limit. The Bandwidth Guarantee allows the traffic of a specific priority to use at least up to the defined percentage even when there are higher priority flows actively running. Bandwidth Limit is the maximum bandwidth allowed for the total traffic of each priority level. All of these parameters can be individually configured for each priority.

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 (blue) 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 were 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 assigned the same priority, the bandwidth would have been equally divided among the three due to the fairness within the same priority.

Flow Classification

AppEx IPEQ currently supports the following classification methods:

- Layer-4 classification
- Layer-7 classification
- Behavior-based classification

Dynamic and Static classification. Aggregations are created as a result of flow classifications.

Layer-4 and Layer-7 classification

As was discussed earlier, layer-4 and layer-7 classifications are available in most of the Flow Control devices today. Although very useful for a 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.

Behavior Based Classification

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 (Bit Torrent, eDonkey, etc.) or a VoIP application (Skype, Live Messenger, etc.) without having to dig into the details of each individual implementation. In many cases, Behavior-Based Classification can be more efficient and effective than layer-7 classification because of its adaptive and intelligent nature. Since Flow Control itself regulates the behavior of the traffic, classifying by the traffic pattern is more accurate than classifying by layer-4/7 content. Behavior-Based Classification is a more viable solution in the long run.

Dynamic and Static Classifications

Dynamic and Static Classifications are concepts independent from the layer-4/7 and behavior-based cllassification examples above, which dictate how the Aggregations are created. A static classification is simply an Aggregation created for a group of flows that match a certain rule or pattern. A dynamic classification is an Aggregation 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 from a set of criteria. 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 which can be deployed easily and at minimal cost. AppEx IPEQ is able to tackle complicated deployment scenarios, offering bidirectional end-to-end Flow Control with minimum packet loss and fairness with prioritization among the flows.

Common Use Case Scenarios

AppEx IPEQ is effective in a wide variety of scenarios thanks to its scaling, bidirectional shaping and its feature-rich Flow Control methods. The major advantage of AppEx IPEQ lies within its inbound traffic shaping. As illustrated earlier, the traditional inbound rate control causes wasted bandwidth. This is the very reason many customers found that some network applications became unbearably slow after the deployment of other Flow Control devices. Very often they were forced to buy more bandwidth from their ISPs, which is an expensive solution. AppEx's inbound traffic shaping solves the problem while using bandwidth more efficiently. This protects the customers' investment and saves bandwidth cost.

Businesses and Organizations

Today's businesses rely heavily on the network to run their operations. At the same time most of them also provide Internet access to their employees at work. Competition for limited

network resources and the high cost to buy more bandwidth is a constant challenge for the IT organization, especially between branch offices.

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

Figure 6: P2P Traffic Control

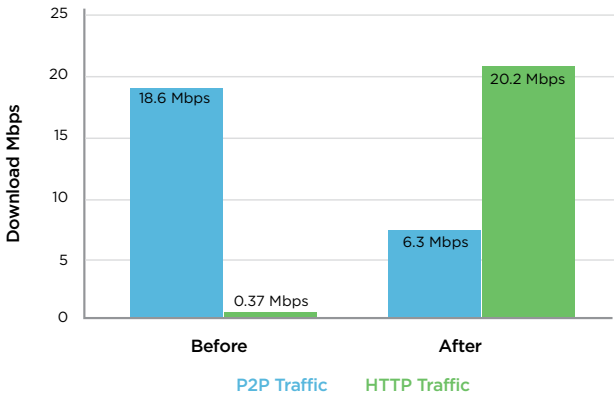

Schools and Public Network Providers

In the schools (especially colleges and universities) and other similar environments, the network applications vary greatly due to all sort of different needs. The IT department usually has little control over what applications the end users utilize on the network. Very often P2P-based applications occupy most of the bandwidth, leaving the majority of the normal network applications to suffer. The same situations also occur in certain public businesses, such as Internet Cafes. The P2P applications use a wide variety of protocols, and in the academic community with creative and inquisitive minds, there is a great deal of bandwidth demand. In these environments AppEx's behavior-based classification with host fairness approaches are well suited.

Figure 6 shows the bandwidth usage before and after the deployment of AppEx IPEQ for a customer. The bandwidth subscription from the ISP was 30Mbps. Before using AppEx IPEQ, the P2P application was able to utilize 18.6Mbps while HTTP was able to utilize 0.37Mbps. In the same scenario with IPEQ deployed, P2P traffic was limited to 20% of the total bandwidth or about 6.3Mbps while HTTP was able to increase throughput to 20.2Mbps. This was accomplished using Behavior-Based-Classification.

Figure 7: Bandwidth Prioritization

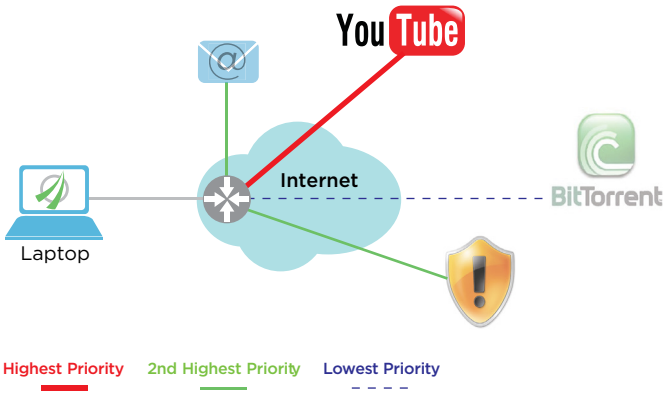

Home Use

Most homes now own multiple computers and network capable devices. As a result bandwidth competition within the home is fierce. In the majority of the scenarios the preferred solution would be to have all the devices share the bandwidth fairly and prevent any single application from taking over all the bandwidth. There might also be certain applications that home users wish to prioritize over others, such as VoIP calls, online video streaming, online gaming, etc.

IPEQ's small footprint allows it to easily be built into home routers, which have limited hardware resources. AppEx behavior based classification is capable of identifying P2P, VoIP, video and other traffic flows and categorizing them into Aggregations. With AppEx host fairness and a few other predefined Flow Control rules the router will be able to keep everyone happy.

Personal Computers, Tablets and Smart Devices

Until recently it has not been necessary for PCs, Tablets, or Smart Phones to perform any kind of advanced Flow Control. With the rapid development of high-quality media streaming, network gaming, VoIP and live communications, etc., the demand for desktop Flow Control has increased. It is gradually becoming an issue because bandwidth heavy applications are competing for the bandwidth.

Figure 8: AppEx Application Bandwidth Controller

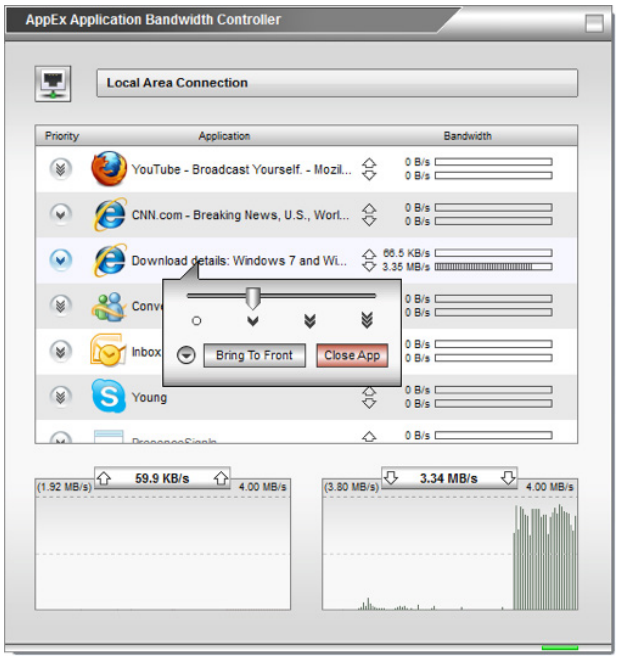

AppEx IPEQ (Application Bandwidth Controller) is automatically configured when the software is loaded on any PC, Tablet or Smart Phone. The software automatically redefines dynamic host-Aggregations into process-Aggregations. This redefinition enables the Application Fairness feature so that no processes will be able to deprive others of network access by hogging the bandwidth. This feature is on by default.

IPEQ also enables users to prioritize traffic with a simple user interface shown below. This UI lets a user manage the priority of each application in real time. For example, online games, VoIP and online video applications can be marked high-priority to ensure the best performance for these applications.

Conclusion

Bandwidth competition is the major factor affecting 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 is the next generation Flow Control technology that is able to accurately control inbound traffic, minimizes packet loss and wasted bandwidth, and eliminates congestion on last mile. The combination of the bidirectional traffic shaping, prioritization and fairness control enables the best bandwidth efficiency and quality of user experience.

AppEx Networks Corporation
1601 McCarthy Blvd.
Milpitas, CA, 95035
+1 408-973-7898

More information can be found at:
www.appexnetworks.com

LotClient software has IPEQ functionality.
For a Free LotClient trial:
download.appexnetworks.com

