Info-Communications Development Authority of Singapore

Net Neutrality Consultation Paper

Anagran Response

December 27th, 2010

TABLE OF CONTENTS

I. EXECUTIVE SUMMARY

- Problem Statement
- Anagran's Fair Use Policy
- Addressing Net Neutrality Constituent Concerns
- Response to IDA's Net Neutrality Policy Position

II. THE BATTLE FOR CONTROL OF THE INTERNET

III. THE PROPOSED ANAGRAN FAIR USE MODEL FRAMEWORK

ADDENDUM – FLOW RATE CONTROL: OPTIMIZING NETWORK BANDWIDTH UTILIZATION IN A FAIR AND BALANCED MANNER

I. EXECUTIVE SUMMARY

Problem Statement

An increasingly contentious debate between Consumer groups, Network Operators, and Internet-based Content/ Service Providers is raging in the industry over “net neutrality”. The debate has essentially condensed into two camps: those who argue that any management of broadband network traffic will severely harm the Internet and technological innovation; and those who claim the right to impose traffic management rules and pricing structures that they deem necessary to achieve an acceptable return on investment.

In the meantime, IP network congestion on broadband networks continues to grow as new, bandwidth-intensive services and multimedia applications evolve and proliferate. With its current legacy infrastructure and not originally designed for the delivery of quality sensitive applications such as voice and video, today’s router-based Internet is incapable of handling the associated network traffic volume without negatively impacting the performance and quality of experience of business and consumer users.

QoS DIMENSION	Voice	Video	Data
Bandwidth is too low	**Gaps** *“What did you say?”*	**Pixilation** *Replay / repeat visual*	**Inefficiency** *Retransmit data*
Delay is too high	**Interactivity** *“Are you still there?”*	**Interactivity** *Resynchronize*	**Data lost** *Reconnect and try again*
Jitter is too high	**Gaps** *“What did you say?”*	**Pixilation** *Replay / repeat visual*	**Inefficiency** *Retransmit data*
Loss is too high	**Call lost** *“Please hang up and try again.”*	**Picture lost** *Restart visual*	**Data lost** *Reconnect and try again*

Regulators and Policy Makers seek balance where Internet users have access to content made available on the Internet while ISPs, telecom network operators and Internet companies and content providers maintain the flexibility to a) develop new business and service delivery methods to meet the needs of the market and b) deploy network management techniques that maintain reasonable Internet access quality of service for users, executed in a competitive yet non-discriminatory manner.

Anagran's Fair Use Policy

Anagran offers a Fair Use Policy Framework and its advanced Flow-Rate technology as a proposed solution to the Net Neutrality debate based on its ability to employ noninvasive dynamic IP Flow State information to identify and manage Internet and Multimedia traffic in a way that ensures equal access to online content and facilitates tiered service offerings.

Anagran believes that the Fair Use Policy Framework provides a practical and effective way to address the legitimate concerns of all parties in the Net Neutrality debate, delivering per user and IP traffic protection **without bandwidth** and **performance infringements** (See attached technical white paper on 'Flow Rate Control'). By managing network traffic in a way that is rational and fair, the Framework obviates the need draconian measures such as traffic blocking. And by using noninvasive technologies to identify different types of Internet-based traffic, the Framework also addresses the legitimate concerns regarding user privacy.

The Anagran Fair Use Policy Framework allows applications to be identified as traffic flows using industry-standard information contained in a packet's header without any need to examine the packets' payload. Detecting the type of application without requiring invasion into the actual data payload is essential to establishing rational priorities (e.g. among Emergency Calls vs. other VoIP traffic, file transfers vs. web browsing, email vs. Internet gaming, etc.) and managing bandwidth in a fair way.

Once the myriad applications are identified by their traffic flows, a hierarchy of 3 management policies can be applied:

1. **Proportion Bandwidth Equally Among All Active Users:** Anagran's Fair Use Policy Framework proportions and guarantee minimum bandwidth among all active users. This addresses the problem of a few users of high-bandwidth applications 'hijacking' bandwidth from other average users, and thereby maintains a quality Internet and Multimedia experience for all users. Any unused bandwidth will be fairly and dynamically distributed to the active users.

2. **Prioritize Traffic for Regulatory Compliance:** The Framework prioritizes available bandwidth for critical traffic, such as emergency calls or compliance with Law Enforcement Act. This prioritization can be extended to calls between first responders, or other traffic essential to the public safety,homeland security and online medical consultation services.

3. **Provision Available Bandwidth into Nondiscriminatory QoS Tiers:** The Framework allows network operators to provision available bandwidth into quality of service (QoS) tiers, which may include different rate structures based on guaranteed bandwidth consistently at all times. These QoS tiers treat all traffic within a given tier the same, and so are inherently equitable.

Addressing Net Neutrality Constituent Concerns

Anagran crafted the Fair Use Policy Framework to address the legitimate concerns of the four main constituents in the Net Neutrality debate.

a) For Consumers, it provides nondiscriminatory access to all Internet sites(other than banned sites by MDA), services and content, and eliminates the inherent unfairness of one size fits all rates that now force most users to subsidize high-bandwidth users.

b) For Internet-based Content and Service Providers, it offers different levels of QoS to accommodate the various multimedia applications with a nondiscriminatory allocation of appropriate bandwidth guarantee.

c) For Network Operators, it assures an adequate return on the investment required to expand overall capacity, and provides a means other than controversial traffic filtering and port blocking to enhance QoS by managing available bandwidth more efficiently.

d) Finally, for Regulators and Policy Makers, the Framework establishes a reliable method for prioritizing traffic that ensures public safety and fairly balances the needs of all concerned.

Response to IDA's Net Neutrality Policy Position

Anagran concurs with IDA's policy position regarding net neutrality, specifically a policy where Internet users have access to all legitimate content made available on the Internet while ISPs, telecom network operators and Internet companies and content providers maintain the flexibility to a) develop new business models and service delivery methods to meet the needs of the market and, b) deploy network management techniques that maintain reasonable Internet access quality of service for users, executed in a competitive yet non-discriminatory manner.

Anagran's Fair Use Policy Framework supports IDA's three-pronged approach to facilitate a competitive Internet access market and safeguard consumer interests, while at the same time providing sufficient flexibility for ISPs or telecom network operators to differentiate their business models, services and products, and perform network management. The Fair Use Policy Framework:

- Enables the delivery of tiered services in a fair, equitable and competitive manner;
- Reduce any incentives for ISPs and telecom network operators to engage in blocking or discriminatory conduct that restricts consumer choice; and focus on user broadband plan differentiation based on ISP's contention ratio and subscribed International Bandwidth;
- Enables information transparency for consumers, to include clear visibility to network management activities and the consumer's Internet usage experience, and;
- Protects consumer interests by ensuring that competition in the market does not lead to ISPs or telecom network operators degrading the Internet access service quality to end-users by insuring both compliance with QoS and the deliver of expected performance.

II. THE BATTLE FOR CONTROL OF THE INTERNET

In the current debate over net neutrality, the issues have been cast as the ultimate battle for control of the Internet. The concept of net neutrality, ensuring that consumers of broadband services can access desired Internet sites, services and applications while also ensuring that Internet-based providers are not inhibited in their ability to offer content, services and applications - is not new and has been debated in policy circles for years. However, the debate over net neutrality has grown increasingly vitriolic and strident. It appears that various segments of the industry have segregated into two warring camps, each proposing laws and policies that are mutually exclusive and leave little room for reasoned compromise.

The real issues of maximizing consumer choice and unrestricted access to the Internet, responsible management of network bandwidth allocation in the face of increasing congestion, elimination of barriers to entry for Internet-based businesses, and the appropriate methods for recovery of network costs, are complex and deserve a more responsible public hearing.

To Some Extent, Everybody Has a Point

The constituents in the Net Neutrality debate fall roughly into four groups:

- Business and Consumers Users;
- Internet-based Content and Services Providers;
- Network Operators; and
- Regulators/Policy Makers.

A. Users of Internet-based Content and Services. Users are represented by a number of think tanks and advocacy groups that are active in this debate. These groups reflect public anxiety over several issues and desire, including:

- Assurance that users will not be prevented from accessing the content and services of their choosing, nor forced to use services or providers they do not want;
- Protection against unchecked increases in service rates over time;
- Assurance that the quality of their Internet access will not be degraded, and;
- Protection of personal privacy.

Implications: Gradual migration from centralized to distributed information models

- Migration from centralized content (Video, Audio, etc.) architectures to a distributed Peer to Peer content distribution models (e.g. Google, Facebook, Rapidshare)

- Migration from a centralized computing / storage model to a grid computing based model (e.g. Oracle / IBM / HP upcoming business models)

- Increasing need for reactive network / traffic engineering models, need to revisit the fundamental maximum utility based resources to network allocation schemes (traffic control, routing decisions, decoupling of identification and location)

<u>Constituent Network Recommendation: Upgrade IP Networks for Multimedia Services</u>

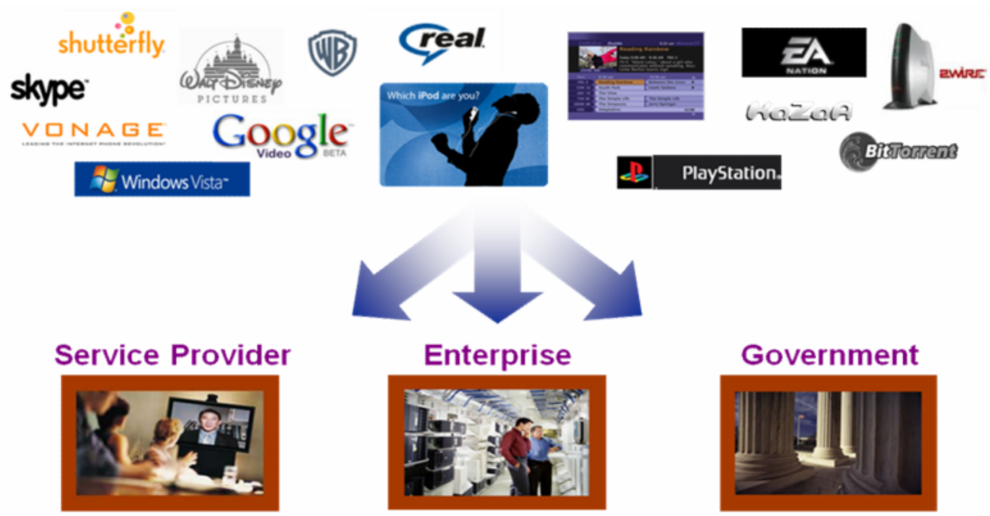

B. Internet -based Content and Service Providers. The most active players in this constituency include Yahoo!, Google, Facebook, Rapidshare, Twitter and trade and industry groups. This group takes the fundamental position that the Internet has developed into the driver for technological and economic development that it is today because it has been unregulated and unrestricted.

Implications: Emergence of overlay service providers (e.g. Skype, etc.) – Disruptive competitive landscape

- Residential and Enterprise market: inflection point in service take up, Security / QoS issues solved – clear impact on existing revenue streams

- Similar shift in wireless as in wireline, leveraging upcoming wireless data broadband technologies (Wifi Mesh, Wibro, 3G/LTE, etc.) – Most significant hit on existing revenue streams

Constituent Network Recommendation: Demand for High Quality IP Multimedia Services

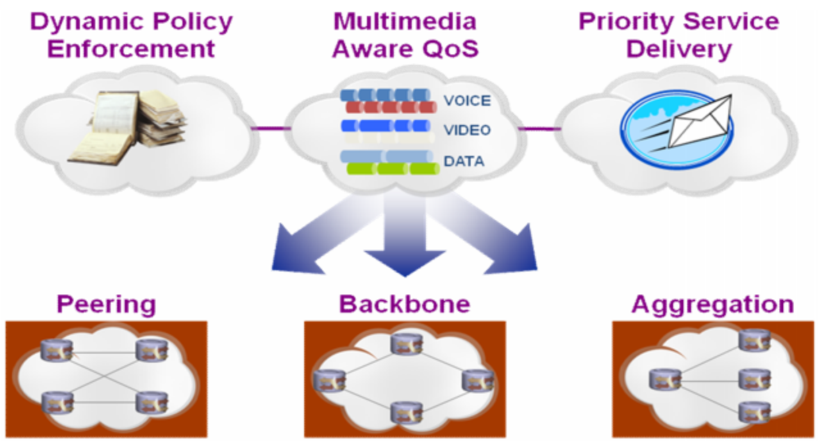

C. Network Operators. This constituency is comprised of the incumbent telephone, wireless and cable companies that provide broadband Internet access and backbone services to users. The network operators argue that use of the Internet is growing explosively to the popularity of high-bandwidth multimedia applications. A small number of high-bandwidth users are causing a disproportionate amount of congestion, and that network operators need the bandwidth management tools to maintain performance for the majority of users. The network operators contend that they should be able to impose charges on, or share revenues with, providers of Internet-based content and services because those providers use the broadband networks in order to reach their customers.

Implications: Competitive landscape shifting from being geography specific to geography agnostic (e.g. Network Operator competes with Google, Yahoo, Facebook) with non-Telecom players increasing business competition

- **Solution #1: Leverage regulatory models to protect business – Risky in the long term**
- **Solution #2: Provide a competitive offering locally and/or globally – requires shift in business practices**
- **Solution #3: Decouple the network from the service, be highly competitive in either or both**

Constituent Recommendation: Enhanced Dynamic Policy Enforcement

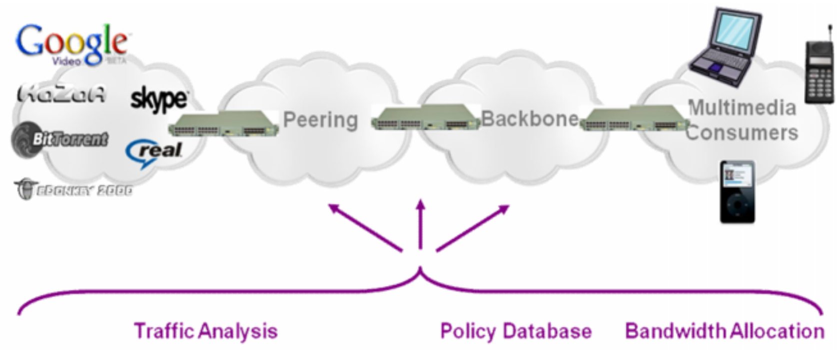

Service	Duration	Rate	Byte Count	Length	Address
HTTP	Short	High	Variable	Short	Unknown
VPN	Long	Low	High	Short	Known
Games	Long	Low	High	Large	Unknown
Streaming	Long	Medium	High	Large	Unknown
Telnet	Long	Low	Medium	Short	Unknown
P-2-P	Long	High	High	Large	Unknown

D. Regulators/Policy Makers. Regulators and other policy makers seek a balance in competing positions, while addressing some fundamental public interest goals. This constituency has its own set of priorities that include:

- Maintain a balanced market in which content/service providers and network operators are free to compete without regulatory impediments or advantages
- Maximize customer choice among Internet-based network operators and content/service providers; and
- Providing absolute priority for critical public needs, including privacy protection, reliable access to emergency services, and meeting the needs of law enforcement and homeland security agencies.

Implications: Regulations

- Multi providers model for residential broadband
- Focus on new revenue streams, with shift to FTTH not likely to bring direct revenue streams and shift to next generation wireless broadband access not likely to increase ARPU
- Global Service Providers will gradually increase competitive pressures on service offering
- Strategy focus on bundling ASP and broadband access services – challenges

Application Matrix

 Multimedia Ready	Dynamic Policy Enforcement	Multimedia Aware QoS	Priority Service Delivery
Service Provider	Fair-Use Enforcement	Triple Play Delivery	Premium IP Services
Enterprise	Application Assurance	Multimedia Services	Mission-Critical IP
Government	Public Access	eGov Services	Netcentric Warfare

Neither of the Polarized Positions Will Work

The respective positions of all four constituencies raise some reasonable and legitimate concerns. But the overly simplistic and fully polarized policy positions advocated by some fail to address the collective concerns in a comprehensive way. One can segment the various positions into two general positions:

- Unlimited Net Neutrality
- Unrestricted Managed Access.

Unlimited Net Neutrality: The most ardent proponents net neutrality describe the transformative effect the Internet has had on an economy and culture, and its technological leadership in the world. These parties ultimately take the position that the introduction of bandwidth management or quality of service tiers will be the end of the Internet as we know it. It is at this point that the rhetoric gets in the way of reasonable policy making.

The proponents of unlimited net neutrality generally ignore the fundamental reality that increased Internet usage has resulted in increased congestion. The fact that Internet usage is growing at a rate of 100% annually and the fact that bandwidth-intensive multimedia applications are growing exponentially are generally ignored.

Instead, many of the proponents of unlimited Net Neutrality argue that future developments, such as new compression technologies, the growth of massive bandwidth delivery with fiber to the home, and broadband network overbuilding by third parties are all factors that might eliminate congestion problems. Of course, the advocates of this position are not the parties who will be asked to make the substantial investments needed to expand network capacity sufficiently to eliminate congestion, momentarily.

Unrestricted Managed Access: On the other side of the fence is the absolutist position in favor of no restrictions on network operators to manage Internet access and impose charges. These proponents generally ignore concerns that network operators may be able to abuse their position in order to collect excessive fees, or gain an unfair competitive advantage against other content or service providers. As with the unlimited net neutrality proponents, the unrestricted managed access will address these concerns through increased competition based upon new services that utilize broadband over power line, satellite, wireless and other technologies, thereby providing a guarantee against unreasonable conduct by network operators.

While these alternatives may provide some competition at some point in the future, they do not offer competitive access alternatives for the vast amount of users today, nor will they address the issue of congestion of the access layer of the network.

III. THE PROPOSED ANAGRAN FAIR USE MODEL FRAMEWORK – A LEVEL PLAYING FIELD FRAMEWORK

Anagran believes that the polarized positions discussed above represent a false dichotomy that unreasonably constrains the choices available to the industry. So, what is “fair use”? Ideally, business and residential users should have access to legitimate content with a reasonable quality of experience and, in a non-discriminatory manner, network providers should be able to allocate capacity and be paid according to what they carry over the network.

We should aim for making networks the best - and most profitable and prolific—that they can be. A fair use policy balances the desire to enable any customer to connect to the provider’s network while ensuring customers who require premium network services, such as QoS and high bandwidth, receive them. This also includes ensuring that abusive applications do not absorb capacity without control by allocating bandwidth based on what users pay.

Network providers can begin implementing fair use by adopting uniform usage policies that clearly state their positions on traffic management. For example, all traffic will be carried, but it will be managed with the goal of ensuring fair treatment. Providers can partner with content providers, developing a uniform policy for bandwidth usage and revenue sharing. They can partner with existing service provider customers, establishing prioritization policies that assure various traffic types are handled uniformly. Now, instead of losing money to access providers, long-haul providers can compete on the quality - instead of simply low price. As recent events in other industries demonstrate, nobody wins in a race to the bottom of the market.

With a balanced fair use policy, the network providers can be appropriately compensated for delivering prioritized traffic relative to bandwidth and QoS features consumed. Consumers benefit because they receive their content faster with a high-quality purchase experience. Content providers can differentiate themselves with faster service and incentives for customers to purchase additional services. A win-win-win for the industry constituents.

The Anagran Difference

While traditional packet-based networks have advanced in networking efficiency, their fundamental strength is also their main limitation—routers and switches only look at individual packets. They can’t see a communication—such as a voice call or a music selection—in its entirety. Unable to see the entire communication, or ***flow***, it is impossible to assign the correct quality of service (QoS), priority, or other attributes to it, as well as intelligently manage all of the other traffic contending for bandwidth at the same time.

The quantum step forward to realizing the potential of IP networking requires shifting the focus from managing individual packets to managing individual communication ***flows***.

Anagran's Flow State technology recognizes individual flows, routing the first packet of the flow and dynamically associating state with it. All remaining packets in the flow are switched using this state information. In addition, for each flow Anagran captures and maintains forwarding, QoS, and real-time statistical data. Flow State management achieves granular control over each individual flow, which provides network managers with the accurate data they need for creating and delivering network services with per flow guaranteed QoS.

Flow-State QoS assigns and maintains QoS data for each flow, including flow priority, class, rate guarantees, connection admission control enabled, delay, variation and jitter. Before a flow is set up, Anagran determines dynamically how much network resources are needed to meet the specific requirements of the flow.

Anagran's technology results in a Fair Use Policy Framework that provides what the label promises: a flexible framework for establishing traffic management policies that are fair to all who use the broadband service.

Business Benefits and Drivers

Anagran's Fair Use Policy Brings the Following Benefits to the Net Neutrality Debate

- **Real-Time User Differentiation, Guaranteed Bandwidth and Flow Protection (Broadband: Fixed Line, 3G/LTE, WiFi/WiMax; Enterprise: Departments, Server Farm, etc.)**

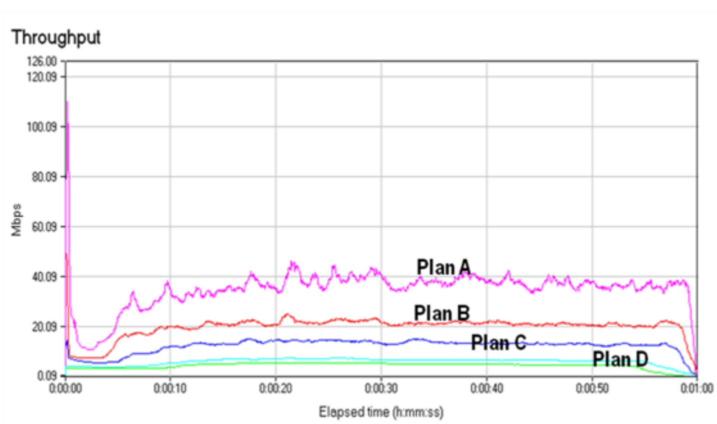

- **Mobility Broadband Congestion Management (3G/LTE, WiFi/WiMax, NFC, Satellite, etc.)**

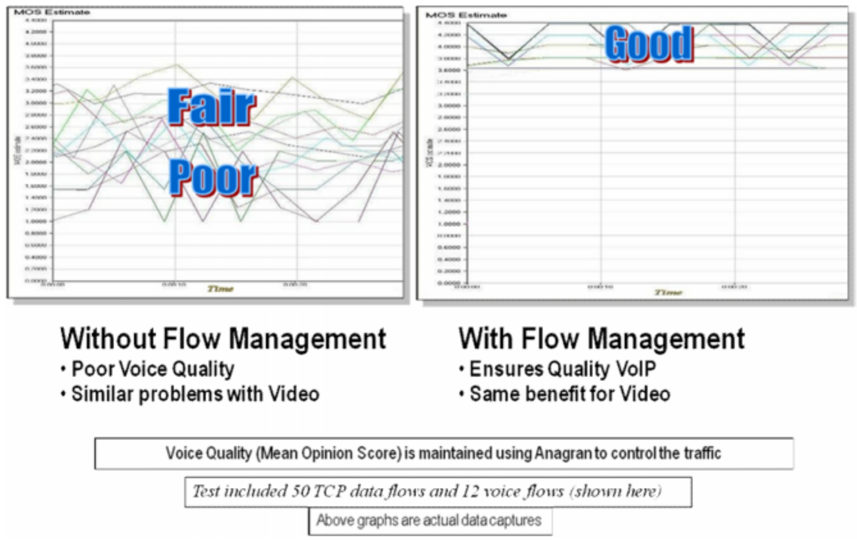

- **Maximum Utility of Network Investment – 95% Bandwidth Utilisation**

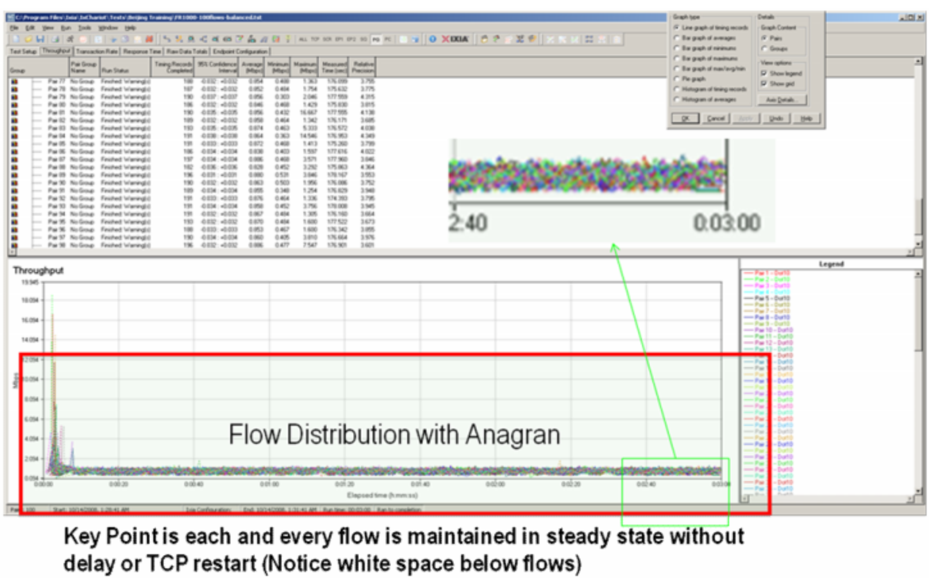

- **Quality IP Video with Dynamic Policy and Delivery Management**

- **Predictable experience for users all the time**

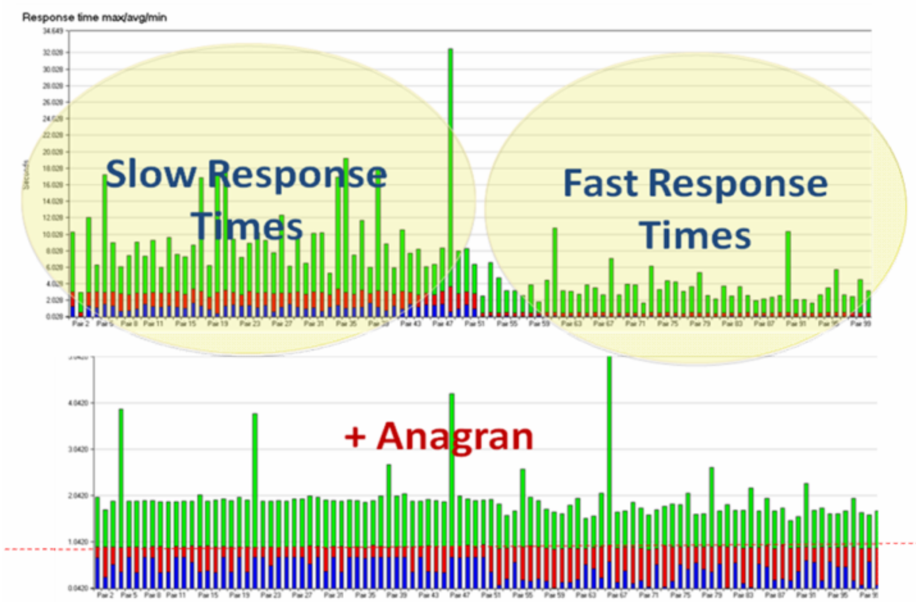

- **Congestion Avoidance**
  - Maximise Bandwidth Utilisation up to 90 - 95% (rather than from 60 - 70%)
  - Minimal if not zero symptoms of congestion

- **Differentiated User Experience between Policy Groups for Usage Profiling**
  - Special Events/Activities gets guaranteed bandwidth as required
  - Administration and departments gets double/triple the data rate of other normal users
  - Normal User share equally and fairly the remaining bandwidth

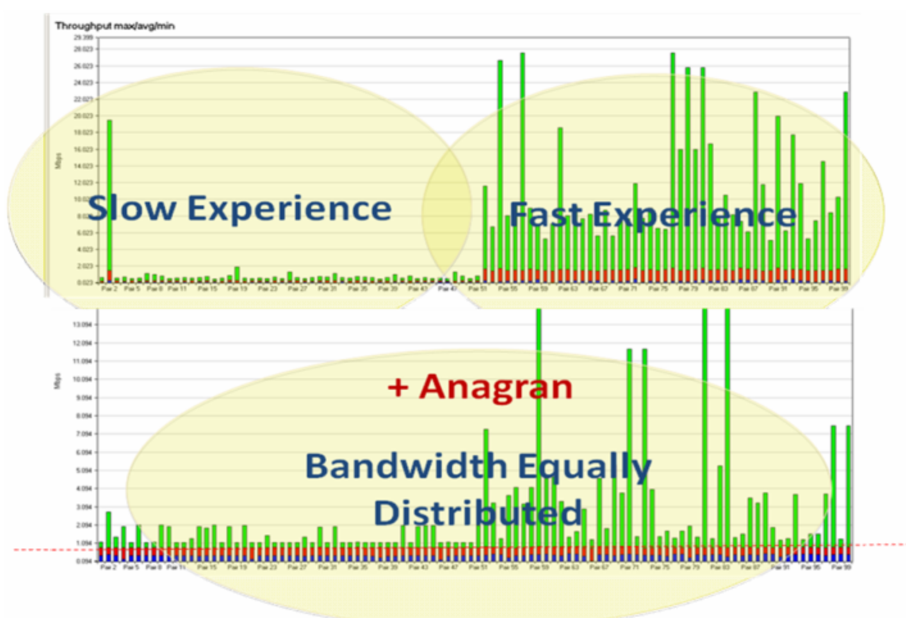

- **Preference and Priority for Certain Traffic**
  - Increased computer usage for all users in enterprise with quality user experience
  - Guaranteed allocation for online traffic
  - Guaranteed allocation for video multicast/broadcast (distance learning, consultation, etc.)

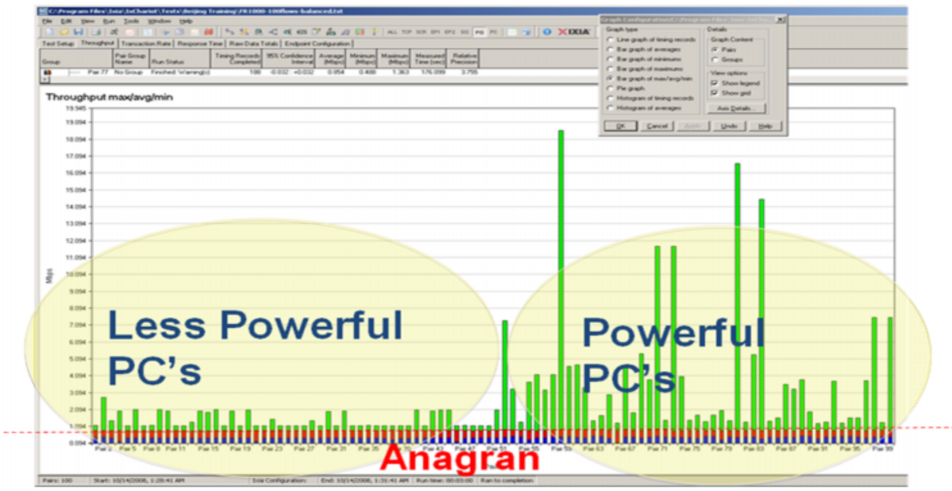

In addition, the Fair Use Framework imposes no requirement for any intrusive means to identify traffic flows and/or applications, such as deep packet inspection, that could compromise personal privacy. For all constituents this means,

- **Easier way to manage IP Services**

- *100% Real Time Visibility …* **For all traffic; including encrypted, compressed, IPv6 etc.**
- *Cost Effective Visibility…***Negates need to support Netflow on other devices**
- *Accurate Visibility…..***Netflow records based on <u>complete flows</u>, not packet samples**
- *Scalable Visibility……***Only platform able to produce complete Netflow records at *<u>up to 10Gbps</u>***

Anagran’s Fair Use Policy Framework Satisfies the Core Concerns of All Constituents

Based on its successful implementation of the Fair Use Policy Framework in Anagran Flow State Management deployed in other countries, Anagran is confident that the Framework effectively addresses the legitimate concerns of all constituents in the Net Neutrality debate.

Consumers of Internet-based Content and Services. A poll conducted by Consumer Reports generally show that consumers primary concerns are that:

- They have continued, uninterrupted access to their favorite Websites and service providers;
- They do not want to be forced to take service from an ISP that is not of their choosing; and
- They want to be protected against rate hikes for Internet access services.

The Framework provides ample guarantees on all three of these fronts. Perhaps the most significant is the ability of optional premium QoS tiers to eliminate upward pressure on the existing one size fits all pricing model that is now forcing some users to subsidize the excessive usage of others. In addition, the Framework eliminates any need for network operators to engage in port blocking or traffic filtering to control congestion, and protects personal privacy by not imposing a requirement for deep packet inspection.

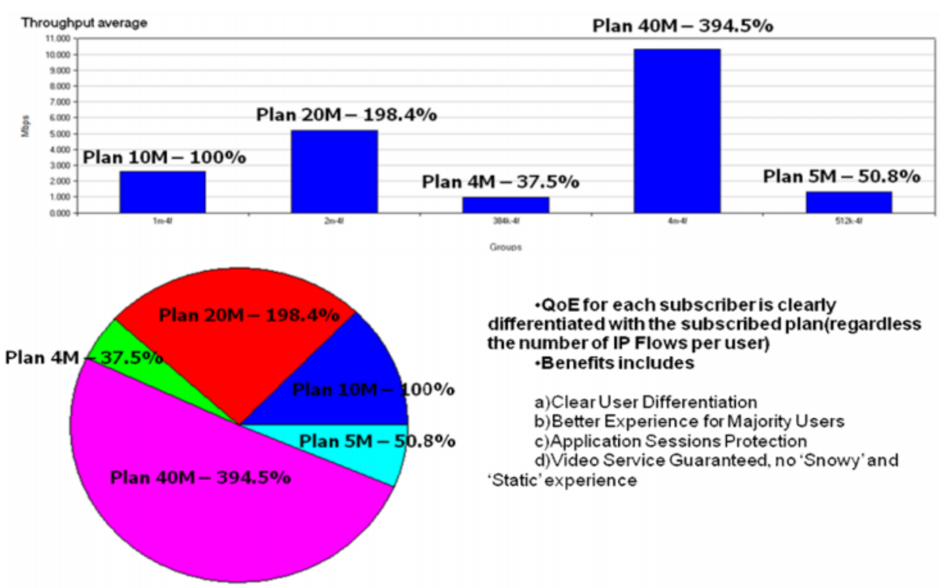

Internet-based Content and Service Providers. Internet-based content and service providers are virtually unanimous in their concern that network operators may provide them lower -quality connectivity than network operators provide their own affiliates, or their preferred content suppliers. The Framework addresses this concern by enabling policies that provide nondiscriminatory access to bandwidth for similar types of applications or services. Content and service providers differ substantially, however, on the issue of sharing revenues or paying network operators for provisioning premium services. Just as many consumers are willing to pay more for more bandwidth, so too are many content/service providers willing to pay more for better QoS for their customers.

Regardless of the outcome, the Framework supports the provisioning of premium services in a nondiscriminatory fashion.

Network Operators. The Framework addresses the three central concerns of network operators:

- A small number of high-bandwidth users and applications cause a disproportionate amount of congestion on the network;
- Additional revenue is needed to recover the large investments being made in network infrastructure. Proportioning allows network operators to dynamically allocate bandwidth equally among all users, effectively preventing a few high-bandwidth users from hijacking available bandwidth from other, average users. Prioritization enables network operators to comply with government regulations, particularly those that pertain to public safety. And provisioning supports the establishment of QoS tiers, which will appeal to high bandwidth users and can generate new revenue streams from the existing customers, as well as from Internet -based content and service providers.
- The Framework allows network operators to manage traffic effectively without engaging in intrusive technologies that could expose them to potential claims of invasion of privacy or misuse of personal information. The Framework also allows effective bandwidth management without engaging in any form of traffic filtering or port blocking, which may be unlawful.

Regulators/Policy Makers. The Fair Use Policy Framework provides the single most effective means for regulators to resolve the issues involved in the Net Neutrality debate while ensuring the priority handling of traffic flows and applications that are in the public interest. The Framework also addresses regulator concerns about the use of intrusive methods for identifying traffic flows, and the use of traffic filtering or port blocking as means of congestion control.

Mobile IP

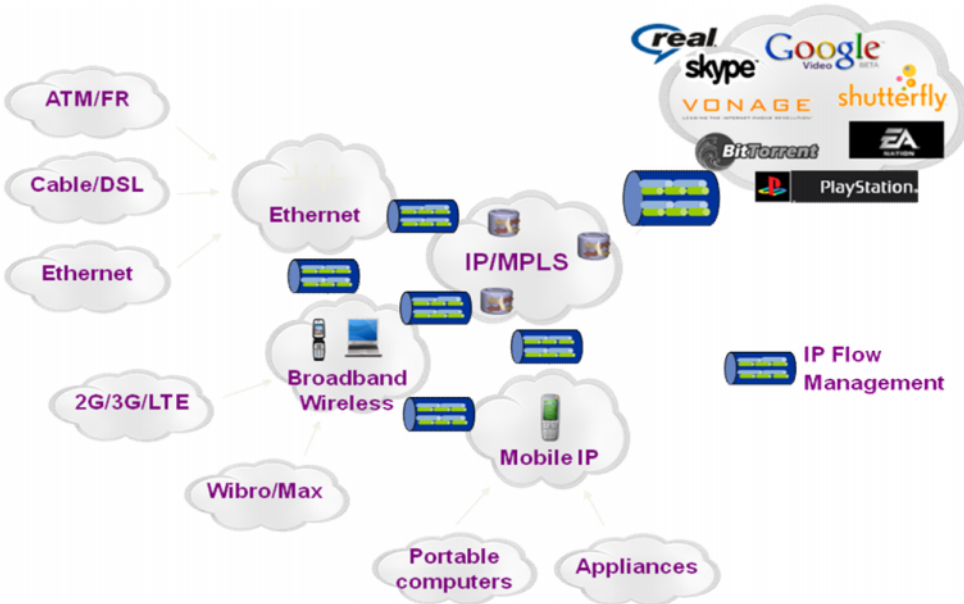

Finally, the Framework provides a robust and versatile means for adopting and enforcing traffic management policies that are nondiscriminatory and, therefore, capable of achieving consensus within the industry.

CONCLUSION

Anagran believes that a versatile and non-intrusive means of managing network bandwidth represents a fruitful middle ground in the increasingly contentious debate over Net Neutrality. The Anagran Fair Use Policy Framework provides a comprehensive three level hierarchy for effectively mitigating congestion and its inherently unfair consequences in a best effort network.

Indeed, it would be difficult to imagine a solution to the problem without the ability to proportion, prioritize and provision bandwidth, or to attempt any form of congestion control without visibility into traffic flows and the nature of the applications involved. It would also be difficult to imagine a solution that required intrusive analysis of traffic, or controversial measures like traffic filtering or port blocking.

For these reasons, Anagran believes that a bandwidth management framework is necessary, and that the Fair Use Policy Framework proposed here is sufficient.

To learn more about the Anagran Fair Use Policy Framework, visit Anagran on the Web at www.anagran.com. Or contact Anagran's representation for the Asia Pacific region.

ADDENDUM

FLOW RATE CONTROL

OPTIMIZING NETWORK BANDWIDTH UTILIZATION IN A FAIR AND BALANCED MANNER

Anagran FRC Network Bandwidth Manager

Optimizing Network Bandwidth Utilization

Internet Traffic and the Bandwidth Management Challenge

The exponential growth in global Internet traffic volume is being fueled by continued broadband service deployments, enterprises turning towards cloud computing services, the rapid growth of smartphones and other intelligent mobile devices with always-on Internet connections, and the change in traffic mix from email and other bulk data applications to rich media and time-sensitive services. TeleGeography's Global Internet Geography study reveals that international Internet traffic grew 62% in 2010 after realizing 74% growth recorded in 2009. Furthermore, it has been forecasted that global Internet traffic will increase more than four-fold by 2014. To accommodate such increased traffic demand, significant network capacity has been added by broadband service providers. In 2010 carriers added 13.2Tbps of new international capacity, up from 9.4Tbps in 2009, and 6Tbps in 2008.

A phenomenon and problem that just about every carrier or institution worldwide has seen is that although no matter how much capacity is added to their network, it is used almost instantly as it is the nature of Internet Protocol (TCP/IP) to consume all the bandwidth that is available. Capacity added one day could easily be fully consumed the next day. Furthermore, although the traffic volume growth is high, typically a very small percentage of users consume the vast majority of capacity on any given network. Most Broadband Service Providers (BSP's) have found that as little as 4-5% of their users are consuming up to 80% of all their network capacity, leaving a limited amount of capacity for the vast majority of subscribers. In an environment where growth of quality sensitive applications such as video, voice, and gaming are accelerating at exponential rates with online video alone (Internet, TV, VOD and P2P) forecasted to exceed 90% of all consumer traffic by 2014, the negative implications to service quality and profitability are huge.

Thus the challenge for BSP's, be they wireline, wireless or cable network operators, is how to optimally improve **network utilization** and improve the users' quality of experience (QoE) in the most efficient, cost effective manner.

This white paper explains how advances in TCP/IP traffic management, particularly the inclusion of per-flow rate control network bandwidth elements in today's Internet network infrastructure enables BSP's to significantly increase network capacity utilization and improve the QoE for all users.

The Underlying Technology Problem: TCP Behavior and the Impact of Queues Used for Network Congestion Control

TCP/IP packet networks, including the Internet, operate as best effort networks with no assured QoS, typically displaying considerable delay, delay jitter and major bandwidth rate variances especially at network congestions points. These attributes are caused by the interaction of TCP with queue based network elements. The problems are aggravated for the rapidly growing Real Time Traffic (video and voice) being carried by wireless and wireline packet networks.

As packet traffic traverses the Internet, TCP grows the data rate, sending larger and larger blocks of packets in bursts until a packet is reported lost by the receiver, at which point the TCP sender drops to half rate and then starts to grow again. This causes (1) the rate to increase until the path overloads and packets are randomly discarded during periods of network congestion or (2) all the flows in the path grow until they have reached their "maximum rate" when the network is not congested.

The current generation of routers and other network equipment use queues to avoid the overload of the next network link or trunk. The queues are simple buffers which drop packets if the queue is full. In addition, Weighted Random Early Discard (WRED) is currently used where a small percentage of packets are discarded if the queue is partially full, and all packets are discarded if the queue is full. All such methods create multiple impacts on TCP with several undesirable effects of handling network congestion.

Delay: Because TCP rate is dependent on the Round Trip Time (RTT) of the path, one immediate impact is the delay caused by the queue which slows down TCP flows. Typically, bottlenecks in any path occur at one or both edges of the network. Thus, if a TCP flow crosses two edges where overload tends to occur, then the delay in the two network edge queues add to the round trip transit time, potentially in both directions. Often the queuing delay is greater than the transit time, adding 40 ms at each edge queue for a total of up to 200 ms.

Delay Jitter**:** Perhaps the most harmful aspect of the added queuing delay is that it varies as the load varies. Thus in the prior example, the delay could vary anywhere from 40 ms to 240 ms, creating a 200 ms delay jitter. From a technology perspective, such large delay jitter is unacceptable for real time applications such as VoIP and video. Given the desire to improve the users' quality of experience, coupled with the increased revenue opportunity associated with the delivery of high quality VoIP and video services, there is a substantial incentive to eliminate the jitter.

Stalled TCP: When a queue gets full and must drop all the new packets, some TCP flows suffer from losing more than one packet. Since TCP bursts a group of packets each RTT, the flows' whose bursts just arrived at the queue generally lose two, three or more packets from one burst. TCP rules then tell the sender to drop to slow start (one packet per RTT). Even more harmful are the cases where the TCP process sees the loss of more packets in one RTT and goes into a stall. The stall is a period with no transmission and this period is extended longer and longer if additional losses are detected.

The ramifications of such stalls can be significant. For instance, a typical web access may have 40 flows, each delivering 20-60 kb of data. In a network with even moderate overload, these stalls will occur on many of the 40 flows. The slowest flow determines the display time for the web page. Queues spread out the distribution of flow rates so much that web pages are typically three times slower to complete than would be true if the flows kept about the same rate. In fact, experiments have shown that if the stalls are eliminated one can improve web page response time by a factor of 3X.

Synchronization**:** When the TCP traffic on a trunk approaches 50% utilization it tends to synchronize. Initially a few flows peak at once and thus the queue drops some of their packets. If the flows' Round Trip Delays are similar they will tend to peak again at the same time. When this occurs additional flows get caught in the peak and the aggregate of all flows peak together. At some point all the TCP flows with similar RTT's will peak together and the channel will suffer considerable delay jitter and packet loss. This is why most trunks need to upgraded when they exceed 45% utilization in the peak hour. However, if the packet drops are organized to avoid synchronization utilization can be increased to about 95%, thus doubling the effective capacity of the trunk.

The Solution: Flow Rate Control (FRC) Network Bandwidth Management Elements

To briefly summarize the balance of this paper, flow rate control (FRC) systems accurately manage the rate of every flow with the composite rate of all flows in any subset of the traffic to be maintained within 90% - 95% of the capacity of a multiplicity of network congestion points. Three critical benefits are realized: 1) significant improvement in the users' quality of their Internet experience, 2) material improvement in network capacity utilization, lowering both Service Provider's CAPEX and OPEX, and 3) increased revenue generation opportunities through the delivery of tiered services.

Flow Rate Control: From a macro perspective, a flow rate control-enabled network element looks into a flow table and finds a match for the five IP packet fields that define a flow (source address, destination address, protocol, source port and destination port). These IP packet fields determine where and how the new packet should be sent on the same route computed for the first packet of the flow.

The flow rates, as well as other behavioral parameters, are collected that allow the TCP flow rates to be precisely controlled at the input to a system. If the flow rate is excessive, one and only one packet is dropped such that the TCP flow operates at precisely the rate that the flow should be allowed. Clearly if all flows could be managed at an aggregation point, the total rate of any group of flows could be controlled such that when that group arrived at a congestion point there would be no need for delaying packets or for dropping packets randomly at the congestion point. Figure 1 shows the structure of a flow rate control system.

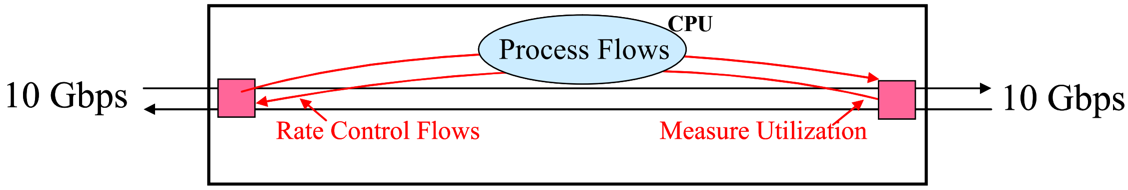

Figure 1. Flow Rate Control System

Flow Rate Priority: The historical concept of capacity allocation for a group of flows was to divide capacity equally between all flows. This is the natural outcome of TCP and queues and was acceptable when the Internet got started since typically each user used one flow and this made all users equal. If one flow was sending more packets than another, the random discard at the queue would on the average discard more packets from the larger flow, slowing it down such that all flows tend to the same rate. Today, however, a more flexible allocation of rates per flow, flow rate prioritization, is required to allow each flow or subset of flows to receive different rates. Flow rate priorities can be used so that the rates received are proportional to the priorities.

Flow rate control systems also provide benefits to User Datagram Protocol (UDP) and layer 2 protocols. By precisely controlling TCP rates, non-TCP traffic can be protected and prioritized, minimizing the effects of delay jitter and reducing the pressure on higher latency TCP queues. This improves the experience and responsiveness of streaming traffic, VPN sessions, or other quality-sensitive traffic. Alternatively, non-TCP traffic can be controlled just as TCP can, giving an equitable amount of bandwidth to each flow. This protects other flows from aggressive applications that use UDP or tunnels attempting to circumvent TCP back off mechanisms and use more bandwidth than they would normally receive. FRC's are also capable of transparently performing UDP-based Call Admission Control (CAC), only accepting new sessions if sufficient bandwidth exists to support it and existing sessions without reducing their quality.

Impact of Flow Rate Control on Network Congestion at the Broadband Edge

In normal usage of the Internet the primary point where sufficient overload occurs is at the broadband edge between the BSP and the customer. For enterprise customers, traffic management will be most effective at the corporate edge given the worst congestion point is generally the connection to the Internet. For residential broadband customers, the congestion point is usually at the BSP's service concentration points for subscribers, i.e. the DSLAM, CMTS or GGSN backhaul links.

Net Neutral Management of Abusive Bandwidth Users

One of the major problems facing BSP's today is that a small fraction of the users consume the majority of the pool of network bandwidth capacity intended for a entire community. To date, P2P applications have been the primary offenders. By its very nature a P2P application opens many flows to download music or videos. Each flow tends to receive equal bandwidth due to TCP's interaction with queue-based congestion management. P2P applications having 100 flows receive 100 times the capacity of other single flow applications. As application designers have determined that the ability to open several flows increases overall application throughput, other non-P2P applications are starting to utilize multiple flows.

The FCC in the U.S. has ruled that Service Providers should not discriminate against a particular application but can exercise fair and equitable traffic management to maintain good service levels. Anagran's patent-pending FRC technology is uniquely suitable for providing fair bandwidth allocation in an application agnostic method. Anagran FRC network bandwidth managers measure all the traffic and flows per subscriber and in cases of network congestion, rate control all flows such that multi-flow users receive the same capacity as the other normal subscribers. This is referred to as "subscriber equalization". FRC is not fooled by encryption but is application agnostic thus meeting the FCC requirements for a net neutral method of bandwidth management.

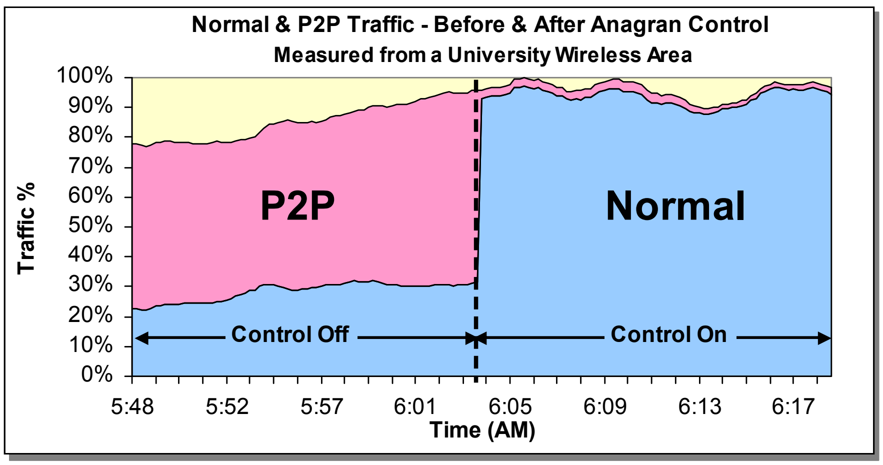

Figure 2: Example of Subscriber Equalization

A specific example of FRC control of P2P traffic overloading a network is shown in Figure 2. P2P bandwidth hogging is the most aggravating problem caused by the "equal capacity per flow" rule. In measurements made at a University which installed Anagran FRC network bandwidth managers, P2P dominates the capacity until subscriber equalization was turned on at 6:03AM. At that point the P2P traffic received its fair share of bandwidth and the average user's capacity increased fourfold.

The solution for Service Providers is to deploy FRC technology, implementing "equal capacity for equal pay". Where Service Providers have deployed Anagran FRC network bandwidth managers for subscriber equalization, Quality of Service (QoS) and Quality of Experience (QoE) radically improves for all users, with multi-flow users still receiving their fair share of capacity.

Prioritizing Subscriber Traffic

One extremely valuable feature of the Anagran FRC network bandwidth manager is the ability to assign a rate priority to flows, subscribers and/or traffic classes.

Subscriber Priority*:* Depending on the service a given subscriber pays for (e.g., gold, silver, or bronze), Anagran's FRC system can set the designated flow rate priority for each subscriber.

For example, assume gold customer gets a rate priority of "4", silver "2", and bronze "1". The gold user will have a flow rate four times the rate of the bronze subscriber's flows and twice the silver subscriber. Each class of subscribers can also be assigned a traffic rate cap which controls their peak rate during non-congestion periods. With both these controls a BSP can provide an extremely effective boost in their network performance and productivity.

Traffic Class Priority: One way to get much improved utilization from a fixed size Internet access link is to establish traffic classes, say bulk traffic and interactive traffic, and then prioritize such traffic accordingly. Anagran's FRC network bandwidth manager can identify bulk traffic with a "behavioral" command. For instance, the command would say: "If the number of bytes received on a flow exceeds 100,000, it is reclassified as bulk". Then if "bulk" is assigned a priority of ½ of interactive, the bulk traffic will flow rapidly if there is little interactive traffic but be suppressed when the interactive traffic peaks. This lets the bulk traffic fill the valleys between the interactive peaks rather than pushing the peak up even higher.

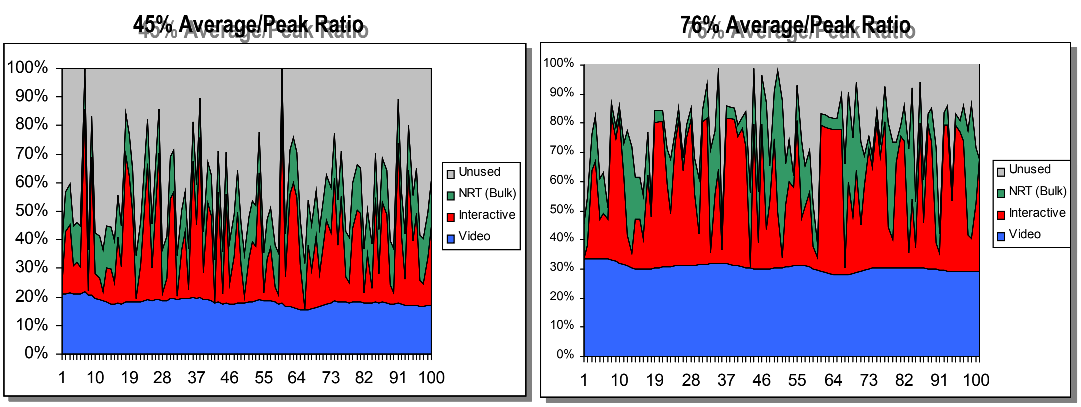

Figure 3: Bulk Traffic without (left) and with (right) Traffic Class Prioritization

Figure 3 is an example which in the left table indicates equal traffic class priorities as compared to the table on the right where interactive traffic has twice the priority of bulk traffic. The response time for interactive traffic is the same in both cases but the throughput of all three types of traffic has been increased 68%.

Summary

A fundamental challenge for network operators is the tremendous increase in wireline and wireless Internet traffic volume and the consequential effect on overall network performance, CAPEX and OPEX investments, and opportunities to increase ARPU via tiered services. As traffic is increasingly generated by applications sensitive to delay and jitter (voice and video), any effect even of temporary congestion on a network link results in a very noticeable decrease in the quality of the service by the end user.

Anagran's Flow Rate Control network bandwidth managers can be deployed transparently in the existing network infrastructure consisting of legacy packet-based networking devices to provide a significant improvement in network efficiency and quality of experience. Rather than managing individual packets with their associated increased delay and jitter, flow rate control technology manages flows per subscriber or traffic class. With a consistently predictable service quality, network operators have a greater opportunity to build profitable services for their customer base without the necessity of expensive capacity increases. With Anagran's unique flow rate control technology this approach immediately yields Service Providers improvements in the efficiency of link utilization, lowering CAPEX and OPEX, improving ARPU and their customers' overall user experience.