

Net Neutrality QoS

Win More Subscribers with Smarter QoS
In ISP environment

PacketController White Paper

support@packetcontroller.com
www.packetcontroller.com

March 2011

Contents

Executive Summary	3
The issue of DPI Methodology.....	5
Per Subscriber Methodology.....	6
Process of QoS Implementation.....	7
Strategy.....	7
Implementation.....	8
Integration.....	10
Technology.....	12
Conclusion.....	14

Executive Summary

Bandwidth management for a service provider is different from a normal businesses requirement because it directly effects profitability and can be a primary service differentiator for marketing and sales purposes. To stay competitive, service providers must be able to support customer satisfaction, cost reductions and cross-sell/up-sell efforts. To do that, however, means that service providers must have scalable, reliable and flexible QoS solution in place. Better yet, they should be guided through all necessary steps from strategy definition to QoS policy to maximize on ROI between the backhaul investment and subscriber services.

Traditional DPI (Deep Packet Inspection) solutions fail to address the key issue of bandwidth management in ISP: subscriber volume; DPI solution require many QoS policies for application recognition. However QoS policies based on explicitly having to identify the application are problematic as there is always going to be unidentified traffic as signatures change or worse still traffic becomes encrypted. This traffic is then thrown into an "all other" classification and managed in a single umbrella policy. Multiple traffic types some good some bad having to compete for restricted bandwidth. There are many legal forms of p2p downloading as well which get restricted by these general catchall shaping policies.

This white paper describes PacketController's net neutrality methodology, a more efficient and less labor-intensive solution. Through PacketController's automatic TCP rate limiting technology, and it does not matter what the traffic is or if its encrypted or anything as the algorithms at this level fundamentally work per subscriber. So should there be any unidentified traffic (and there always will be) PacketController's unique fair allocation of bandwidth automatically ensures fairness across the network and ensures no one subscriber gets an unfair share of bandwidth. In the end a subscriber using all his bandwidth on encrypted p2p for example will get the same amount of bandwidth as an elderly person trying to do his or her web surfing.

The paper also describes the typical process to implement PacketController in ISP environment to manage bandwidth investment and increase network performance.

Finally the paper presents the features of PacketController's technology, designed and built using the latest technologies and industry standards to provide maximum flexibility and interoperability with existing systems.

Together these capabilities form a comprehensive, easy-to-deploy bandwidth management solution that helps ISP maximize the subscriber satisfaction with the high profit.

The issues of DPI Methodology

Deep packet inspection (DPI) is Application-based traffic optimization, which uses the properties of each network protocol to provide the minimum bandwidth that guarantees acceptable quality. Bulk file transfer applications are given the lowest priority since they are typically non-interactive and long-lived. For example, a one way bulk interactive application such as a file download would be lowest priority, a one-way streaming media like YouTube[®] may be next in priority and an interactive application such as VoIP would have the highest priority. As the network becomes heavily congested this prioritization becomes important as each application is degraded if it is not prioritized.

Internet standards have anticipated that 'differentiated services' would be offered, where applications 'mark' themselves into the appropriate class based on the priority need of their packets. For example, VoIP marks itself as a high priority given its real-time bandwidth need and a file download marks itself at a lower priority. This provides priority for real-time applications and prevents larger applications from dominating the network. This method, however, is flawed when used in a consumer access application. Broadband access networks (DOCSIS, DSL) do not support 'differentiated services' due to technological limitations. Additionally, differentiated services lead to a fairness issue between subscribers and an incentive to 'cheat', causing the theft of QoS. Application writers sometimes marked their application's packets as the highest priority and this honor system failed.

Service providers have resorted to marking the traffic on behalf of the user, automatically choosing the guarantees that were needed. This application optimization delivers excellent overall quality and subscriber satisfaction.

However, DPI is fundamentally flawed for Internet network services providers:

- I To control user activity it requires many rules and DPI for application recognition. However policies based on explicitly having to identify the application are problematic as there is always going to be unidentified traffic as signatures change or worse still traffic becomes encrypted. This traffic is then thrown into an "all other" classification and managed in a single umbrella rule. It also implies the endless maintenance and application signature upgrade cost.
- I Multiple traffic types some good some bad having to compete for restricted bandwidth. There are many legal forms of p2p downloading as well which get restricted by these general catch all shaping rules.

Per Subscriber Methodology

Internet only provides per-connection fairness for the TCP transport protocol. Bandwidth-greedy applications like P2P that use UDP for bulk data transfer or open many simultaneous TCP connections can easily circumvent this transport capacity fairness and use more than their fair share of the available bandwidth. PacketController can rather easily enforce a per-subscriber bandwidth usage fairness that ensures all subscribers getting on average an about equal share of the available bandwidth, which is particularly important during periods of network congestion.

It solves the fundamental issue in ISP environment: subscriber volume. The advantages of per subscriber methodology are:

- | Heavy users have no negative performance impact on others
- | Fair distribution of available resources among all users
- | No DPI required

Process of QoS Implementation

Strategy

Bandwidth management box is just a tool, you will definitely need an effective strategy to do the bandwidth management right, and bandwidth management box is just to help to implement the strategy.

The following are the items to be considered to define the strategy:

- | Network congestion during the peak time
- | Different service levels for different subscribers
- | CAP subscribers
- | Burst QoS to maximize the subscriber experience and uplink
- | Premium services

	Description	Priority
Network congestion during the peak time	This is the top priority to consider, if there is network congestion then any subscribers will suffer the poor network quality.	Highest
Burst QoS	Burst QoS is so critical for the ISP; you can provide higher quality services to the customers when uplink resource is available to do so. Use it right then you can have more happy subscribers with less bandwidth.	High
Different services levels for different subscribers	You will have different service packages, for example, home broadband package and CIR for enterprise etc., Usage-based charges are still the norm in mobile networks, but even in wireline networks due to the huge disparity in data volumes between normal and heavy users. A possible way out of this dilemma for ISPs and their subscribers is an approach that strikes a balance between flat rates and usage-based charging. The basic idea is to offer customers a choice of which services they require from their ISP – and they are happy to pay for – and which they do not.	High

Premium Services	The premium services are to retain and attract the subscribers willing to pay more on high quality services.	High
CAP subscribers	Many ISP's have different action requirements that are dependant on usage generally. Please note that CAP implies that no network congestion management is triggered by bandwidth management box.	Middle

Noted: Each and every network is different, so is the business model, you need to think about the strategy seriously before the implementation.

Implementation

QoS implementation is straightforward if strategy defined.

The following are the PacketController implementation steps:

- I ***Burst definition based on the uplink bandwidth usage***
For example, the uplink is 500Mbps, and the burst definition is 400Mbps, the service plan 400Kbps could be bursted to 512Kbps for 3 minutes if the uplink bandwidth usage is less than 400Mbps;

While uplink usage is more than 400Mbps, the service plan 400Kbps is just using 400Kbps.

The burst feature is so powerful that you can handle seasonal load dynamically based on the overall bandwidth usage rather than season, which is more accurate; and you can provide higher quality services to the customers when uplink resource is available to do so.

- I ***Assign different service plans to customers with difference service level***
For example, you have 8000 customers, they fall into different service levels like 248Kbps, 512Kbps, 2Mb/s, 10Mb/s or whatever.

Then you could assign different service plans to those 8000 customers by its IP addresses. When it's done, you can forget it and there is definitely no need to look over QoS policies again every time traffic pattern changes.

You don't need to worry about the system performance and PacketController scales well to 45K QoS policies at multiple gigabit speed.

- I ***Punishment on "bad-behaved" customers***
we are not going to pretend that PacketController implementation is the "silver bullet" for "bad-behaved" customers who eat up all the bandwidth assigned to him/her ALL THE TIME.

ISP services are not supposed to be a punishment of its customers; they are supposed to allow you to provide higher quality services to more people at a lower cost.

However you don't want people who download all day;

PacketController provides automatic P2P bandwidth usage throttling technology, it monitors bandwidth usage of each and every online subscriber to identify online 'bad-behaved' subscribers, and there is option of automatic step down of those 'heavy' subscribers in terms of rate limiting or connections.

I ***Subscriber Usage Statistics/Feedback***

To collect the statistics of subscriber bandwidth usage (both real-time and historical) is important, which will help you understand subscriber traffic pattern on your network. The real-time statistics is very powerful tool for you to track down bottleneck on your network. Further, PacketController provides web-based report utility for each and every subscriber to check his/her traffic usage.

I ***Web Cache/Billing Integration***

Web Cache can definitely enhance user response time and thereby user satisfaction. PacketController could seamlessly integrate with External Web Cache System like Squid, it can distinguish between cached and non-cached traffic and manage it transparently. Also PacketController provides database interface for billing system to retrieve bandwidth statistics.

Integration

In most cases, bandwidth management box is deployed in bridge mode, which works as a transparent MAC layer bridge. This implies lots of things to be considered in terms of integration.

The following is the checklist to be considered.

Uplink bandwidth counting consistency

For Internet network service providers, Ethernet technology evolved into Fast Ethernet and Gigabit Ethernet, and continues evolving; in modern networks more and more network operators prefer using Fast/Gigabit Ethernet as a technology for WAN connections.

So what about your uplink ISP counts the bandwidth including both data and Ethernet headers(14 bytes) while your bandwidth management box only counts the data portion? It means the inconsistency in your uplink ISP and your bandwidth management box, which is the disaster for billing.

MPLS

Multi Protocol Label Switching (MPLS) originated from “Tag Switching” a proprietary Cisco development. The technology was originally developed as a mechanism to improve the performance of core Routers. Today those efficiencies gained in core router performance have been negated due to vastly improved hardware technology; however the benefits of MPLS as a service prevail.

Why do organizations elect to implement an MPLS wide area network? In ninety percent of cases it is down to one thing alone, Quality of Service (QoS). MPLS enables the consolidation of applications onto a single network whilst providing the mechanism to prioritize the latency of individual applications within Application Classes. Organizations can optimize their wide area network usage based upon the types of applications communicating across it. The number of application classes varies upon the implementation offered by the service provider but is typically acknowledged as being 3. Each class has a different priority e.g. high priority is for the traffic that requires the lowest latency such as VOIP, medium priority for business critical applications that are not so latency critical and low for those that are unclassified.

Organizations purchase an MPLS service as a base rental cost with supplements proportional to their specified bandwidth for each application class. In return the service provider will provide a performance SLA for each application class.

When deployed bandwidth management box inside the MPLS path, at the very least it should support inspection of IP addresses in MPLS-encapsulated IP packets. This makes bandwidth management in an MPLS path impossible otherwise traffic just goes through not adhering to QoS rules.

Further it will be definitely better if bandwidth management box could add granularity to the bursting process allowing one to choose which applications can dynamically burst in order of priority into remaining unused bandwidth in different classes as Some MPLS providers do support dynamic bursting between classes.

Integration with other software

It includes billing software, monitoring software, Web proxy server etc., it implies that the bandwidth management solution support data interactive with those software, through database, API, SNMP etc.

Conclusion

The conclusion is that the integration is the important factor taking the nature of bandwidth management solution. So it is good idea to get the bottom of integration mechanisms provided by the bandwidth management solutions before purchase.

Technology

PacketController were designed and built using the latest technologies and industry standards, providing maximum flexibility and interoperability with existing architectures.

I ***Congestion Management***

PacketController removes all congestion, drops and retransmission from the network. It does this by effectively pacing traffic of all types to the link speed and unlike the competitors it does this without the need to queue traffic.

PacketController emulates separate links per user, controlling bidirectional throughput to the bit and TCP rate limiting across these to further control congestion. This is the natural way to reduce the number of packets on the network at any given time and therefore reduce congestion allowing higher priority traffic a free passage through the network.

The fastest way for an IP packet to traverse any link is when there is zero congestion between the client and host, by entering the WAN link speed into the PacketController its bandwidth management algorithms will automatically and fairly split bandwidth amongst the users. Through controlling client requests, higher priority traffic always has a free path through the network and never gets delayed in a queue; as a result it paces all traffic to the speed of the WAN link so that there is no congestion at all along the traffic path.

I ***High performance system***

PacketController has the fine-tuned multi-core support for fast parallel packet processing. It uses some well known performance-boosting techniques, such as memory-mapping the card's packet buffers, I/O batching and multiqueue to spread the load across multi-cores. Those technologies allow PacketController to achieve almost straight line performance gains to cores. As of 2011, PacketController 500 could handle 5 Gb/s (Full Duplex) easily on one bridge with 45K QoS policies loaded.

I ***Automatic P2P bandwidth usage throttling***

The automatic P2P bandwidth usage throttling technology brings policy enforcement improvement and user experience enhancement. It monitors bandwidth usage of each and every online subscriber, and it differentiates P2P usage and other applications like web surfing.

Through automatic P2P bandwidth usage throttling technology, the following snapshot information provided:

- * The total number of online 'heavy' subscribers who has been mostly likely using P2P
- * A list of IP addresses of current 'heavy' subscriber
- * The option of automatic step down of those 'heavy' subscribers, for example, service plan switching automatically from 512Kbps to 256Kbps
- * The option of automatic step down of those 'heavy' subscribers through connection per second

I *Integrated With External Cache*

A web cache is a mechanism for the temporary storage (caching) of web documents, such as HTML pages and images, to reduce bandwidth usage, server load, and perceived lag. A web cache stores copies of documents passing through it; subsequent requests may be satisfied from the cache if certain conditions are met.

it is quite common for us to find that an organization already has a firewall and/or proxy in place and we have made sure that PacketController will be compatible with them. PacketController sits behind the Web Proxy, and it is web cache aware which means that PacketController will automatically identify cache and non-cache traffic, this is significant for bandwidth management.

Conclusions

PacketController is used by hundreds of ISP customers worldwide. This demanding customer base has resulted in the highest throughput, lowest latency, most powerful and feature rich bandwidth management product available today.

PacketController provides a comprehensive, out-of-the box solution to address congestion and increased the complexity of network management. Designed to meet the unique and:

- | Optimizes your network usage and performance
- | Moves best effort delivery to a predictable service
- | Dedicates the correct amount of bandwidth to your users and applications
- | increases user productivity satisfaction and performance of all your applications
- | Enforces service level agreements with effective management of your bandwidth
- | Reduces management, support and maintenance costs
- | Allows the business to eliminate and delay costly bandwidth upgrades

PacketController reduces network congestion, discards and retransmissions, increasing performance, reliability and integrity providing comprehensive bandwidth management and control.