



---

## Contents

Introduction	1
Challenges	2
Technologies	2
IPDR	3
Deep Packet Inspection	5
PacketCable Multimedia	5
3GPP Policy Charge Control	8
Sandvine Solution	8
Conclusion	12

## Sandvine Network Policy Control Architecture for Cable

Cable operators face both opportunities and challenges in the market today. To meet these, many are upgrading their networks from DOCSIS 2.0 to DOCSIS 3.0 to enable the offering of higher-speed services to their customers. However, DOCSIS 3.0 creates some hurdles. This whitepaper discusses the challenges that operators face with DOCSIS 3.0 as well as the solutions available including Sandvine's solution for maximizing the opportunity.

## Introduction

Cable operators are presented with both opportunities and challenges in the market for their High Speed Internet service. The bandwidth war between cable operators and other wireline operators continues as the telcos move from deploying DSL to FTTx and offering bandwidth in excess of 100 Mbps. The cable operators have nearly exhausted the capacity of the DOCSIS® 2.0, and are in the process of upgrading their networks to DOCSIS 3.0. DOCSIS 3.0 helps operators address the current network speed limitations by enabling them to increase both the upstream and downstream speeds to 120 Mbps and 160 Mbps respectively.

The cable industry views the upgrade of their network to DOCSIS 3.0 as the foundation for the next generation cable network and the services it will deliver. DOCSIS 3.0 is not just about delivering the Internet faster to the consumer; it is about utilizing an all-IP network for delivering a range of services including video. According to Nielsen Online, in January 2009, about 136 million people watched online video content, up 16% from the same period in 2008.<sup>1</sup> One of the services that DOCSIS 3.0 makes better is the viewing of online video content. Furthermore as cable operators build out their wireless extensions to their cable networks they are in the unique position of being able to offer to their subscribers the ability to pay once for the video content and making it available on any screen (TV, PC, or wireless handset) at any time.

The higher speed service only opens the flood gates for more over-the-top bandwidth hungry applications that will continue to consume the bandwidth faster than it can be added. Bandwidth consumption is growing in excess of 40% compound annual growth rate (CAGR) per subscriber and no reason to believe it will level off.<sup>2</sup> To ensure the maximum quality of experience (QoE) for their customers, cable operators will need to ensure the applications using the network work as intended by ensuring that these applications get the bandwidth they need. They will either need to continue to blindly add more bandwidth to their networks or actively manage their network to minimize the impacts of congestion.

At the end of the day, the revenue has to exceed or match the costs required to provide the bandwidth. Simply adding increasing-offered speeds and charging more is not a sustainable long term strategy as their competition, the telcos with FTTx, can simply match speed upgrade for speed upgrade at a much lower cost once the FTTx is in the ground. Therefore cable operators must start employing a two-pronged approach - usage management and congestion management. Usage management to better monetize the over the top services that cannibalize their legacy services and congestion management to ensure the over the top services work well.

The challenges faced by cable operators is further complicated with the wide spread availability of mobile broadband from the wireless operators. Broadband always-on, on-the-go is becoming the norm and not the exception. With the growing availability of mobile broadband, many consumers are contemplating cutting the cord in manner similar to the cord cutting occurring in the voice business.

The key challenges facing cable operators today are:

- Managing the usage enabled with the higher bandwidth as a result of DOCSIS 3.0
- Congestion management during peak periods to ensure all applications work smoothly over the shared media to maximize the consumer's quality of experience
- Making all the same DOCSIS 3.0 services portable and available on any screen, anywhere, at anytime

Cable operators have the technologies today to help them successfully operate a sustainable, profitable high speed Internet service business but this requires a means to use the full capabilities of their DOCSIS 3.0 network to offer revenue generating services to offset the cost of the network upgrade. DOCSIS 3.0 and 4G network technology provide the backbones for the network transport to make this all become reality.

---

<sup>1</sup> S. Schechner and V Kumar, *Cable Firms Look to Offer TV Programs Online*, <http://online.wsj.com/article/SB123509028580728229.html> (February 20, 2009)

<sup>2</sup> Shaw, Terry D.; Proceedings from Society of Cable Telecommunication Engineers Canadian Summit 2009: *Trends in Bandwidth Utilization*

## Challenges

### ***DOCSIS 3.0***

DOCSIS 3.0 provides cable operators the capabilities to match the current speeds of FTTx. DOCSIS 3.0 builds upon the success of DOCSIS 2.0 by introducing channel bonding technology. This technology provides initial speeds of 160Mbps in the downstream and 120 Mbps in the upstream along with the capability to bond additional channels as required for additional bandwidth. In short, DOCSIS 3.0 provides the foundation for the next generation of cable services.

The introduction of higher speeds enabled with DOCSIS 3.0 creates both new opportunities for service providers as well as new challenges. The opportunities provided by DOCSIS 3.0 include ultra-high speed Internet services, business services over DOCSIS, and IP based video on demand or IPTV. The introduction of ultra-high speed Internet services paves the way for even more consumption with the trend toward higher quality (HD) and more always-on, over-the-top services. With the introduction of these new services, service providers will still be faced with the ever growing consumption of bandwidth that eventually leads to network congestion and the degradation of the applications using the service.

Historically cable operators have offered all-you-can-eat data plans, but with the growing trend towards always on bandwidth hungry applications, many are exploring how to manage the costs of providing high speed Internet. To date, the true cost of not having an enforceable usage management policy in place is not known to the cable operators. Almost all service providers recognize that they cannot afford to continue all-you-can-eat data plans. One approach being considered by many cable operators is to implement some kind of quota or usage management policies whether as a means to better capture some of the value provided by the increased speeds and the services it enables or as a means to contain the costs by encouraging consumers to be conscientious in their usage.

### ***Wireless in the Wings***

In addition to managing the costs of their DOCSIS networks, cable operators are trying to add portability or mobility to expand their service offerings with 4G wireless services whether with WiMAX or LTE. A key component to making the wireless extension an economic success will be leveraging their existing infrastructure and offering services that work equally well on both the DOCSIS 3.0 network and the 4G wireless network. 4G networks by definition have speeds that are much more closely aligned with DOCSIS 2.0 and more importantly use flat all-IP architectures. Both of these enable operators to easily extend their DOCSIS networks with a 4G access network, and to provide many of the same services to their customers regardless of the access network to which they are attached.

## Technologies

With the evolution towards DOCSIS 3.0, a number of technologies have been introduced as solutions in anticipation of the problems associated with widespread adoption of ultra-high speed broadband services based upon DOCSIS 3.0. The cable industry recognized early on the need for better usage metering, QoS management as means of ensuring optimal quality of experience, and traffic management. As part of this, the cable industry introduced in the CableLabs standards the Internet Protocol Detail Record (IPDR) as part of the DOCSIS standards for usage metering and PacketCable Multimedia (PCMM) for dynamic QoS control. The market brought forth other non-standard technologies such as deep packet inspection (DPI) which evolved to become de-facto standard equipment in all DOCSIS networks to provide network management, business intelligence reporting, and QoE reporting. The following descriptions provide an overview of how these technologies function and their increasing limitations in this ever-evolving market.

Table 1 Summary of Technologies

	IPDR	PCMM	DPI	3GPP PCC
DOCSIS Service Flow-based charging	✓			
Service Level Agreement Info	✓			
IP flow based charging			✓	✓
Application based charging			✓	
IP flow based QoS control		✓	✓	✓
Application QoS Control			✓	
Tiered Services			✓	
DOCSIS Compatible	✓	✓	✓	
4G - UMTS/LTE Compatible			✓	✓
4G - WiMAX Compatible			✓	

## IPDR

CableLabs® incorporated IPDR into its DOCSIS 2.0 standards. IPDR was introduced into the CableLabs DOCSIS specifications to provide an alternative to SNMP for usage-based billing and in particular bandwidth-centric usage-based billing. IPDR was designed to overcome many of the well known shortcomings of using SNMP in high speed networks, such as unreliable polling resulting in missed data, counters being limited to 32-bits, counters rolling over faster than they could be polled, and additional CPU processing loads on the CMTS.

IPDR is based upon the classic exporter-collector model in which the CMTS in the network periodically exports a batch of IPDR records to a collector using either a streaming protocol or a store and forward protocol. This simple model provides an efficient and scalable means to collect coarse network analytics that were previously only available by SNMP or other proprietary means.

In the IPDR schema as specified in DOCSIS 2.0, each IPDR record provides, on a per DOCSIS service flow, the running count of the octets transported using the service flow. In addition to the usage metric, each IPDR record also includes an SLA report for the dropped and delayed packets for the service flow.

As shown in Figure 1 and Figure 2, the CableLabs IPDR schema reports usage on a per DOCSIS service flow per device, allowing for measuring aggregate network usage per device. IPDR protocol provides a reliable means for operators to collect network usage data on a per subscriber basis in non-realtime and is ideally suited for simple postpaid usage billing. Figure 2 illustrates using DOCSIS service flows per application to get a more granular usage record.

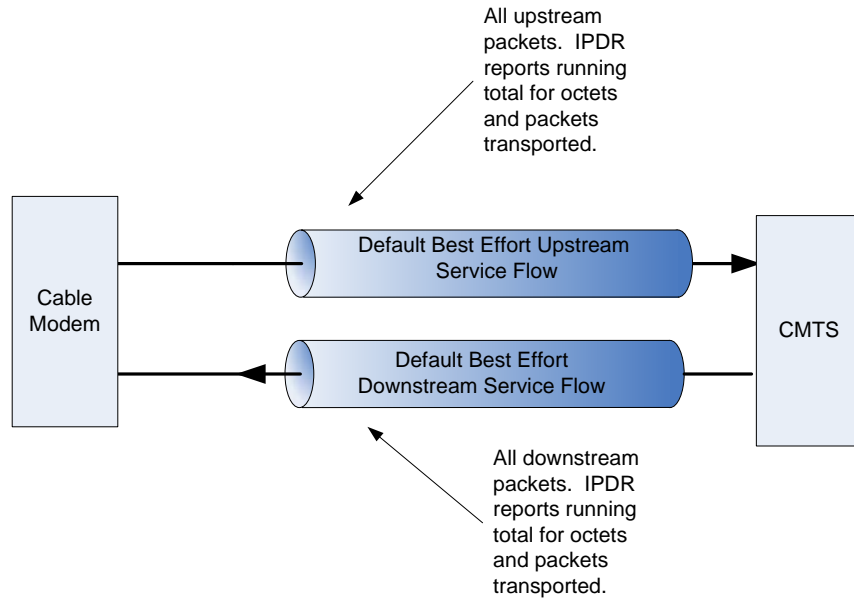


Figure 1 Default DOCSIS Service Flows and IPDR Metrics

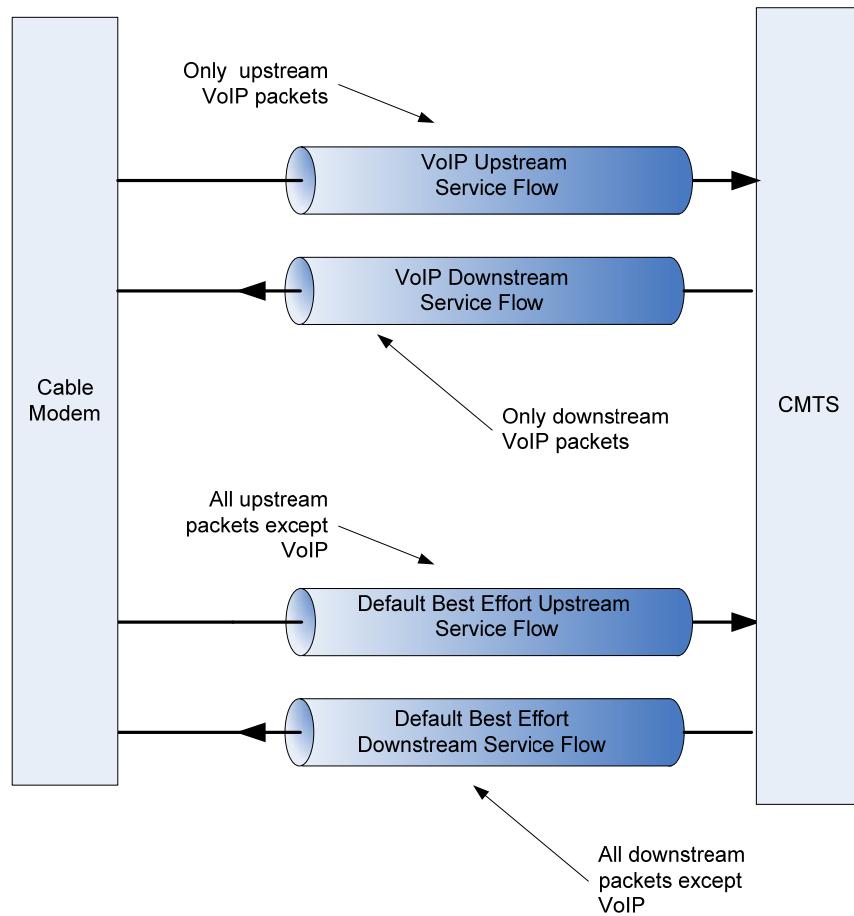


Figure 2 DOCSIS Service Flows with DOCSIS QoS for Voice

The IPDR specification calls for the exporter (CMTS) to transmit a batch of usage records periodically and no faster than every 15 minutes. Therefore, by definition the IPDR data records may be up to 15 minutes. In practice, large and heavily loaded exporters may take even longer to process and export a batch of records. The delayed and non-fixed time grid streaming of IPDR records limits it to non-real time applications. For services that require real-time usage updates such as pre-paid services, quota management, and traffic management, alternatives to IPDR need to be considered.

Further, IPDR does not provide a complete replacement for SNMP in the network management system, as not all the relevant network telemetry is included in IPDR records. This coupled with the fact that fields in the IPDR records such as the SLA information are hard to understand and interpret severely limits the overall usefulness of IPDR and relegates it solely as a technology for non-real time, offline, coarse usage monitoring.

## **DPI**

Unlike IPDR, Deep Packet Inspection (DPI) devices examine the Layer 7 header information of data packets as they pass the inspection point, searching for predefined criteria to decide if the packet can pass or needs some additional processing such as statistical counting, re-routing, or buffering. By examining the Layer 7 header information of the packet, DPI devices provide a much more thorough analysis of the data network. It is an improvement on Stateful Packet Inspection, which provides only inspection via examination of the Layer 3 header portion only.

DPI devices, by definition, have the ability to look at and process layers 2-7 of the OSI model for each packet. This includes the headers and data protocol structures as well as the actual payload of the message. By doing so, DPI devices can classify packets with finer control than classification based upon the 5-tuples in the IP header information with stateful packet inspection. The finer control provides the ability to classify packets on a per application basis and associate these packets and/or flows with the respective subscribers. Once a DPI device has correctly classified the packet, it can then either redirect, mark/tag, block, rate limit, and/or report the packets or flows.

DPI devices operate primarily at layer 3 and above and are thus access network agnostic. This ensures consistent policy enforcement regardless of whether it is a DOCSIS network or some kind of wireless network such as WiMAX, UMTS, LTE, or WiFi. DPI is well suited for metering of usage due to its ability to see deep into the packets. This prevents users from spoofing or cheating the system by running non-standard applications over "well known ports". For example many users will run a bulk file sharing client such as a peer-to-peer client over port 80 to avoid port blocking in firewalls. A well designed DPI device will correctly identify that this is peer-to-peer traffic and not HTTP even though it is using the well known HTTP port, port 80.

However, DPI devices typically are limited by their placement in the network. The placement in the network may affect how much traffic the DPI actually can inspect. The amount of traffic that is routed or injected "below" the DPI device can be minimized by placing the DPI device as close to the edge as possible.

The thorough inspection of the traffic in the network enable DPI elements to generate a rich set of reports ranging from simple network reports that show things like protocol distribution and usage to complex reports that show subscriber quality of experience. Further a well designed DPI system should be able to provide a rich set of business intelligence reports that on how consumers are using the service.

Well designed DPI devices also are usually designed to leverage their power packet classification system for QoS control - shaping, policing, redirect, and blocking. DPI devices work well for downstream shaping and redirection as well as blocking. But, DPI devices are less effective in the upstream for QoS control and as a result need to work closely with the CMTS to leverage the QoS controls in DOCSIS for upstream QoS control.

## PCMM


Unlike DPI and IPDR, PCMM is not a technology for actually metering or enforcing QoS. PCMM is a technology for controlling the QoS on a CMTS. The Cablelabs PacketCable™ Multimedia (PCMM) specification defines a framework for providing QoS-enhanced multimedia services over a DOCSIS 1.1 (or greater) access network. At the core of the PCMM framework is dynamic QoS control in the DOCSIS access network. PCMM is an early form of the Policy and Charging Control (PCC) architecture that has been introduced by the wireless operators as part of the 3GPP Release 7 standards (see the discussion below).

The PCMM standard is intended to address dynamic QoS control, while the 3GPP PCC architecture encompasses both dynamic QoS control AND charging control. The inclusion of charging control in the 3GPP PCC architecture enables flow based charging control in addition to flow based QoS control.

A key concept in PCMM is the PCMM “gate”, with a gate being a logical representation of a policy decision that has been installed on the CMTS device. A PCMM gate on the CMTS performs the traffic classification and enforces the QoS policies on the IP flows which match the gate classifier.

With PCMM, an application can request a dynamic DOCSIS service flow to be created, and have all the packets that match the flow classification mapped to the service flow, resulting in a gate being created on the CMTS. Packets are mapped to service flows using 5-tuple IP classifier (IP source/destination address, source/destination port numbers and IP protocol number such as TCP or UDP) associated with the gate. Each gate has a QoS policy associated with it that defines the QoS parameters of bandwidth, jitter, latency, and traffic priority relative to other DOCSIS flows. In practice, a PCMM gate cannot have a maximum sustained bandwidth that is too small (i.e. < 10 kbps), as this cannot be accurately supported by the CMTS for enforcement.

Figure 3 shows an example of the two static service flows that every cable modem in the network is allocated. The default service flows by definition transport all the traffic. PCMM provides QoS control by dynamically creating adding additional service flows in addition to the static service flows.



*Figure 3 DOCSIS Static Service Flows*

A PCMM gate, by definition, can only add more bandwidth to that already defined for the default service flows. The bandwidth added is not necessarily reserved, but it is added nonetheless. The bandwidth that is added may have a gate that limits the bandwidth to select flows. For example, a PCMM gate could be created that provides 1Mbps of best effort bandwidth for web browsing in the downstream (i.e. traffic with a source port = 80). Therefore, if a user is currently subscribed to a 6Mbps/768Kbps tier of service, he or she would then receive 1Mbps for downstream web browsing traffic (port 80) and 6 Mbps in the downstream for everything else. The net result is that the subscriber has 7Mbps total downstream, but is limited to 1 Mbps for web traffic.

A more complete example is shown in Figure 4. Figure 4 illustrates using PCMM gates to prioritize four different application groups - online gaming, email, web browsing, and video downloads. Each PCMM gate has both a relative priority assigned to it, bandwidth, and a five-tuple classifier. As can be seen in this example the user has actually had 6Mbps of additional downstream bandwidth allocated to them to achieve application prioritization. 6Mbps for the four application categories and 6Mbps for any traffic that does not match one of the five-tuple classifiers assigned to the PCMM gates.

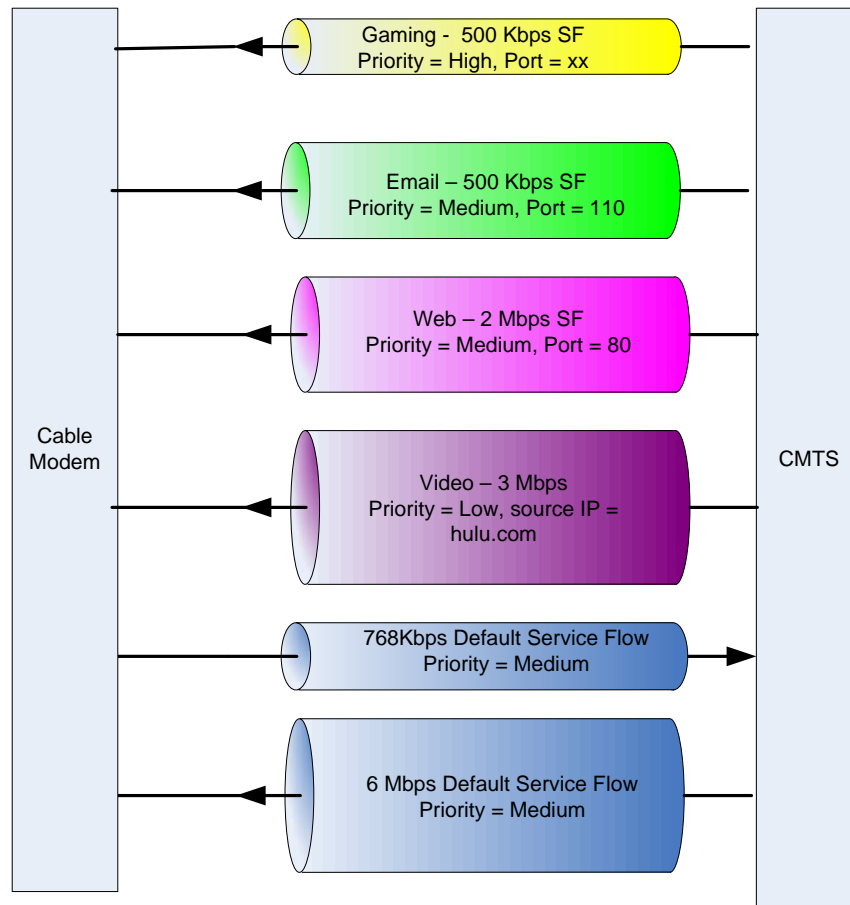


Figure 4 Downstream Application Prioritization using PCMM

As illustrated, PCMM is useful for bandwidth control for IP flows whether to restrict how much bandwidth an IP flow can use or to actually provide more bandwidth for the flow by creating a gate with a maximum sustained bandwidth greater than the bandwidth of the default service flow. In essence each PCMM gate can be thought of as an IP flow bandwidth shaper or policer, with each PCMM gate consuming some router or CMTS resources to enforce the shaper.

However, PCMM has some important limitations that must be considered when using it for QoS control. PCMM gates are a limited resource in both the CMTS and the cable modems (CMs). The gate budgets of most CMTSs are far lower than the number of subscribers supported and therefore it is not practical to design a service requiring one or more concurrent gates per subscriber on a CMTS. Similarly, many cable modems only support four concurrent gates and these gates must be shared with PacketCable 1.x/2.0 services (i.e. digital voice). Furthermore for any flow-based system such as PCMM is that many applications today use many simultaneous flows resulting in large and sometimes dynamic 5-tuple classifiers lists making it impractical to use a flow-based system like PCMM.

Another key limitation of a PCMM gate is that it only can be used effectively for QoS control for an IP flow. As such a PCMM gate cannot be used to block traffic nor can it be used to redirect/re-route traffic. A PCMM gate can be used to mark the DSCP/TOS bit field in the IP header to tell the next hop router to

take some policy based routing action. However, in practice, this is not an effective way to reroute traffic as this is akin to using a hammer when a screwdriver would work. For example, trying to re-route or re-direct all HTTP traffic using this method would result in re-routing all port 80 traffic regardless of whether it was HTTP traffic or not using port 80. Furthermore, it would be easy for users to circumvent this by using one of the many web proxy servers on the Internet.

### ***3GPP Release 7 Policy and Charging Control (PCC)***

Just as the cable industry introduced PCMM, the wireless industry's standards group, 3GPP, introduced the Policy and Charging Control (PCC) architecture in release 7 of its specifications which is also being considered by the WiMAX forum for inclusion in its next release of specifications. The PCC architecture builds upon the PCMM architecture that was introduced by CableLabs for QoS control by extending it to include charging control and to include support for mobility. The PCC architecture provides a standard supports both IP flow based (5-tuple) QoS control and charging control. The architecture provides a standard method for a policy manager to push an IP flow based QoS or charging control policy to an edge access aggregator element such as the GGSN in the 3G architecture or the SAE Gateway in the LTE architecture.

The 3GPP PCC differs from the PCMM in that it specifies the use of the Diameter protocol as the method for the network to export the charging information to the backend systems. Diameter in contrast to IPDR is more flexible in what it can report and the rates that it can support making it well suited for real-time charging applications such as pre-paid data.

Many of the same challenges faced by PCMM are still present because both PCMM and PCC are both IP flow-based architectures and work at only layer 3 and therefore do not allow them to be application aware.

## **Summary**

To summarize, the DOCSIS network has evolved greatly over the last 10+ years since the introduction of DOCSIS 1.0. With this evolution a number of technologies emerged piecemeal to enable the operators to deploy and operate high speed Internet services based upon DOCSIS 3.0. However as the discussion above clearly indicates, no single technology by itself can address all of the challenges that come with the next generation of services enabled with DOCSIS 3.0. Truly addressing the challenges facing the cable operators requires all of the technologies working together in a holistic network policy control framework.

## Sandvine Solution

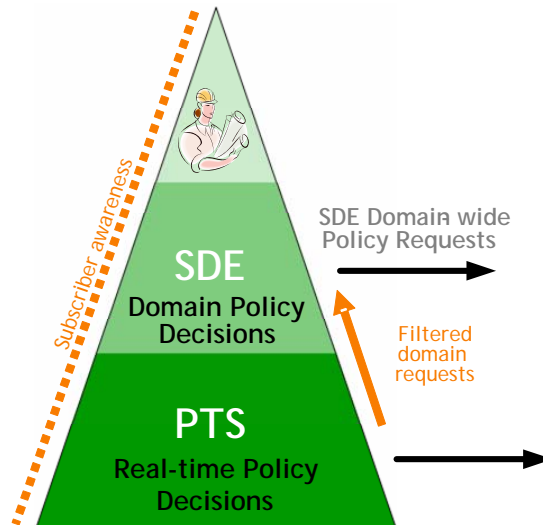


Figure 5 Sandvine Policy Pyramid

To meet these challenges, Sandvine introduced its Network Policy Control Platform. To achieve scale, the architecture uses a two-tier distributed approach for both collection of network analytics/business intelligence and granular policy enforcement, creating a complete policy control solution as shown in Figure 5. The two key elements that are the foundation of the Sandvine's platform are the Policy Traffic Switch (PTS) and the Service Delivery Engine (SDE). Both platforms include real-time policy engines working together, one in the data/bearer plane and the other in the control/domain-level plane. Together, they process the network analytics and events to make policy control decisions.

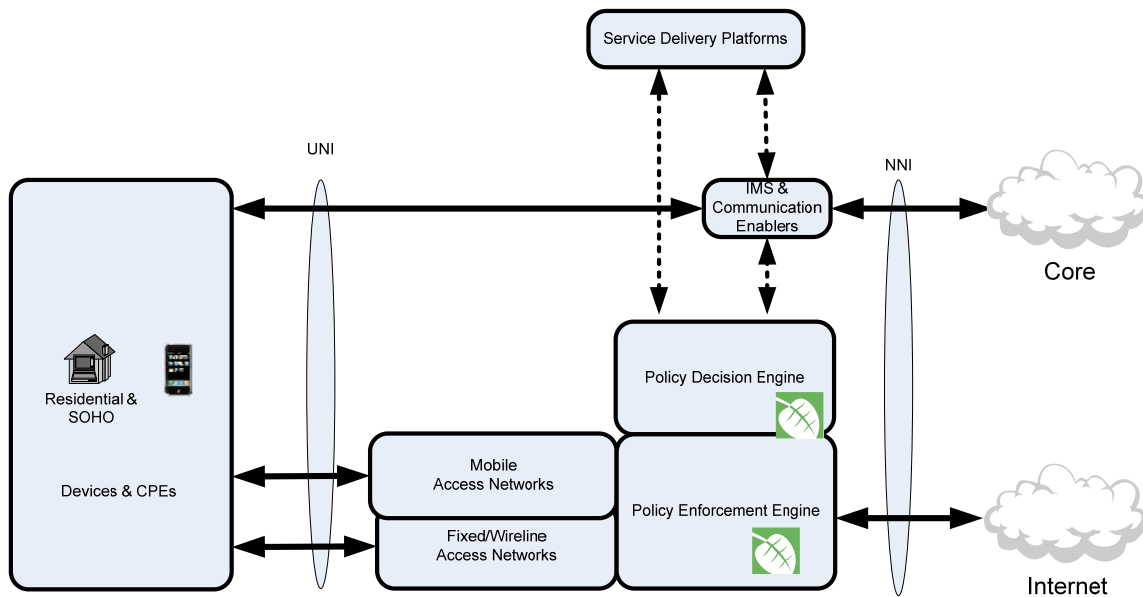


Figure 6 Sandvine Network Policy Control Architecture for Broadband

Sandvine's platform is a holistic, open-standards based solution that provides an integrated policy control architecture that leverages IPDR, PCMM, and DPI for a complete network policy control platform across DOCSIS and wireless networks.

The foundation of Sandvine's platform is the Sandvine Policy Traffic Switch (PTS). The PTS is a carrier class hardware element designed to meet the processing needs of today's largest broadband carriers. The PTS has an embedded real-time policy engine that can process and maintain flow state information for millions of concurrent flows and can make policy decisions in real-time based upon the flow state information for QoS control and charging control.

The key to all of this is the ability to bind the flows to the application, environment, subscriber, and entitlement and network ingress/egress ports, allowing the cable operators to define and enforce flow-based policies as well as application based policies on a per subscriber basis.

The binding of this information together enables the Sandvine solution to produce a rich set of reports that can be used for network engineers for network capacity planning and business intelligence reports and for marketers in their product planning.

Based upon the business intelligence reports cable operators can determine the appropriate subscriber-aware application policies to enforce policies that include measuring usage on a per application basis, enforcing QoS as required, and transparently communicating the policies to subscribers in the right format.

QoS policy enforcement on the PTS can be done two ways; either per flow using a 5-tuple in a manner similar to how QoS is enforced with a CMTS, or at layer 7, the application. QoS policies can be anything from simple shapers as shown in Figure 7, to multi-level priority queuing/policing as shown in Figure 8, to a cascaded, application prioritization on a per subscriber similar to the one shown in Figure 9.




Figure 7 Simple shaping for IP flows or applications




Figure 8 Prioritized queuing per egress port

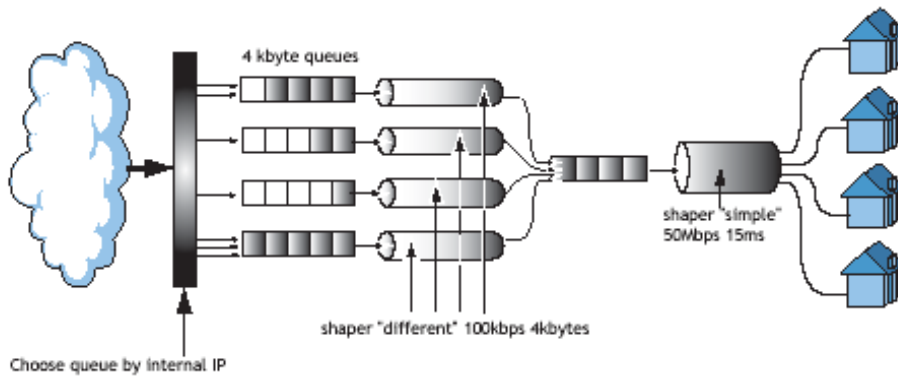


Figure 9 Comprehensive cascaded shaper to prioritize applications per egress port

The second tier of the architecture is the Sandvine Service Delivery Engine (SDE). The SDE includes the logical elements defined in the PCMM architecture -- the PCMM Policy Server and Application Managers -- making the Sandvine platform compliant and compatible with the PCMM specifications. In addition to the logical elements of the PCMM architecture, the SDE includes a rich set of interfaces for gathering of network analytics and application enablers. The SDE includes a DOCSIS IPDR collector for collection of network analytics from third party systems.

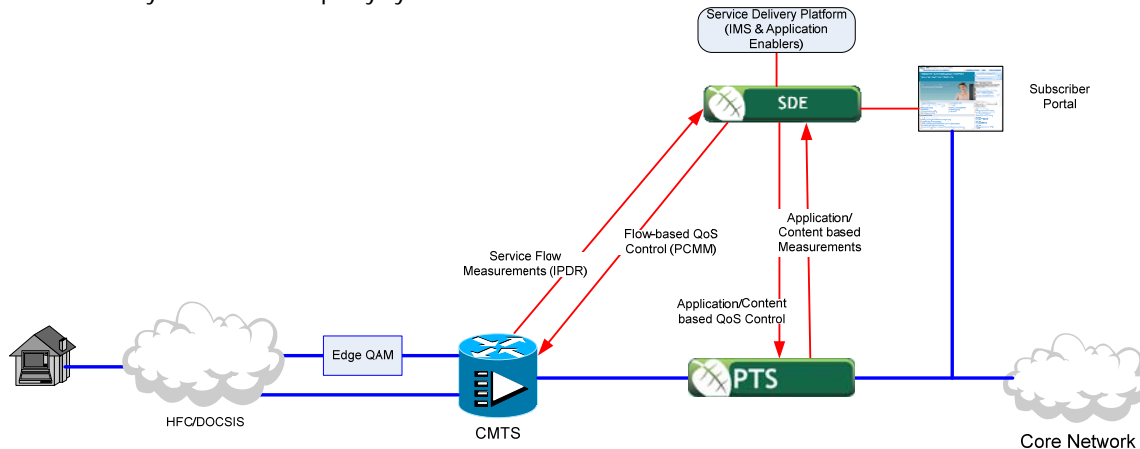


Figure 10 Sandvine Network Policy Control Platform for Cable

The SDE's real-time policy engine processes the network analytics to make real-time policy decisions and signals the policy action to best point in the network for enforcement. For flow-based actions the SDE can trigger policies at the edge of the network on the CMTS using PCMM and for application based actions it signals the policy to the PTS.

The Sandvine platform and its two-tier architecture implement a broad range of policy based services - usage management, fair use, and QoS control. The platform's policy engines work hand-in-hand to process the millions of concurrent flows in a cable operator's network in real-time and then decide the optimal location in the network to enforce the policy. The policy actions can be correlated and overlaid upon the reports to enable the networks to see how effective the policies are in managing the traffic and optimizing the quality of experience for all the subscribers.

Subscriber usage measurement is obtained using both IPDR and the PTS. IPDR data is used for fair use measurements while the PTS is used to measure the monthly usage caps due to its ability measure usage on a per application basis. Doing this on a per application basis enables the cable operators to then zero-rate applications as required including zero-rating software updates and applications like voice.

The fair use policies are enforced at the edge of the network using PCMM. The enforcement of monthly quota caps is achieved using the PTS. For example, the PTS is used to redirect only web/HTTP to a customer portal and not just all the port 80 traffic. This ensures that VoIP traffic using port 80 is not

incorrectly redirected to a web portal. Similarly the PTS may also be used when a usage management policy is triggered, whether it is to block all traffic for an application group, lower its relative priority, or shape it to a lower bandwidth.

The Sandvine platform makes it easy for cable operators to extend the same policy based services to 4G wireless extensions, as shown in Figure 11.

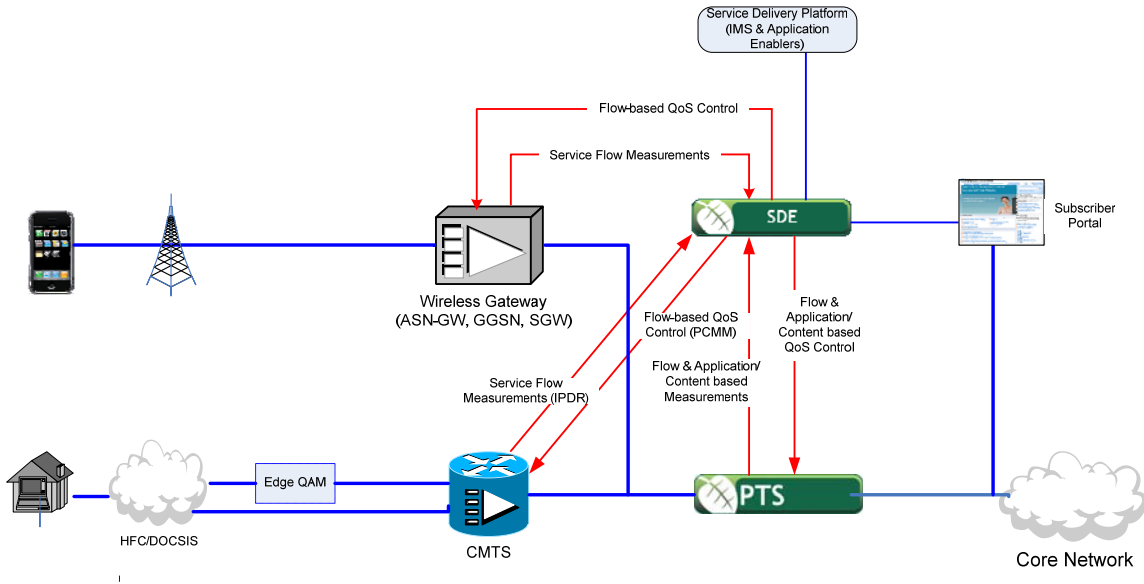



Figure 11 Sandvine Network Policy Control Platform for Cable and 4G

Sandvine’s platform for cable operators eliminates the need for building piecemeal solutions by providing comprehensive, holistic policy control platform. The platform eliminates the need for cable operators to spend time and money integrating individual solutions from multiple operators by providing a complete, working solution today that can be easily extended to 4G wireless as needed in the future.

## Conclusion

As operators roll-out DOCSIS 3.0, three areas for success that they must address are bandwidth usage cost containment and better monetization of offered bandwidth, minimizing the impacts of congestion to maximize quality of experience, and broadband portability. Sandvine’s solution is well suited to address all of these by providing a network agnostic network policy control platform.

Table 2 Summary of Features and Technology

	Sandvine 	IPDR	PCMM	DPI	3GPP PCC
DOCSIS Service Flow-based charging	✓	✓			
Service Level Agreement Info	✓	✓			
Quality of Experience Reporting	✓				
IP flow based charging	✓			✓	✓
Application based charging	✓			✓	
IP flow based QoS control	✓		✓	✓	✓
Application QoS Control	✓			✓	
Application based redirection	✓				
Tiered Services	✓			✓	
DOCSIS Compatible	✓	✓	✓	✓	
4G - UMTS/LTE Compatible	✓			✓	✓
4G - WiMAX Compatible	✓			✓	

The Sandvine platform provides cable operators the ability to define policies that align with their business goals with a holistic system that includes:

- Network/Business Intelligence Reports
  - Subscriber persona reports
  - Network demographics
  - Cross-sectional reports showing the policies and subscriber behavior
  - Quality of Experience (QoE)
- Usage Management Policies
  - Fair use - Policies to address extraordinary use by a small percentage of users
  - Monthly quotas - Total usage quotas as well as budget per applications
  - Zero rating - Identifying applications that are not to be counted against usage caps such as Microsoft updates, cable modem software updates, or online video
- QoS Control
  - Tiering of speed tiers
  - QoS protection for select applications like voice
  - Fair use QoS protection - Shape or lower the traffic priority of the abuse users during periods of congestion
  - Bandwidth on demand
- Subscriber Communications
  - Redirect all the flows for an application to the respective endpoint
  - Communicate in a manner consistent with the application
- Subscriber Personalization
  - Fully transparency to MSO policies
  - Historical reports to when the policies have been applied
  - Customization of policies
    - Budgeting of monthly caps per application category
    - Choice of notification when usage policies trigger or near triggering
    - Choice of policy behavior when usage policies trigger
    - Opt-out option

And with this Sandvine provides the means for operators to leverage their investment in DOCSIS 3.0 to the fullest extent possible and to extent the services to their 4G wireless networks tomorrow.