PER-FLOW QoS ENABLES TRANSACTIONAL BANDWIDTH MANAGEMENT

Gerry White
Chief Technologist
Network Infrastructure Group
Motorola Broadband Communications Sector

Abstract

Transactional bandwidth management holds great promise for the cable industry by allowing improved use of resources through statistical multiplexing and dynamic bandwidth assignment. This must occur in an environment in which Quality of Service (QoS)-enabled applications such as voice place very specific requirements on the network equipment and communication links. Any transactional changes must take place such that these applications continue to function correctly. The price to achieve this is increased complexity in management and control because bandwidth and QoS changes occur in real time. Per- flow queuing can be used to simplify the management and control associated with these changes and provide a deterministic mechanism to implement transactional bandwidth management. This paper explains the importance of per-flow QoS to the successful deployment of transactional bandwidth services and demonstrates how QoS can be provided end-to-end for transaction-based services

INTRODUCTION

As applied in a next generation cable network, transactional bandwidth management is more than simply changing the bandwidth required by a user or application. In many cases, services such as voice, video and streaming audio requires all the attributes of QoS for an application to be changed dynamically. These attributes include not only basic bandwidth but potentially latency, jitter, service interval, and both minimum guaranteed and maximum allowed data rates.

The ability to dynamically apply QoS treatments will allow cable operators and their revenue-sharing partners to offer exciting transactional bandwidth management services where subscribers can automatically provision increased bandwidth for applications such as video conferencing, interactive gaming, or Video on Demand (VOD). Operators can allow the applications themselves to trigger requests for more access bandwidth and application-specific QoS configuration. The self-provisioning of high-margin services is dependent on the ability to automate QoS control across the access network, metropolitan network and the core network of multiple providers.

Bandwidth is and will continue to be a scarce resource, but the Internet Protocol (IP) can be used to support the QoS requirements of multiple services. Operators can ensure that bandwidth is successfully allocated across multiple applications while maintaining the ability to monitor and account for each traffic flow.

Transactional bandwidth management implies that the QoS requirements on the network systems are constantly changing as new applications start, stop or change their QoS needs. In order to manage this potentially chaotic situation it is critical that operators be able to inspect individual traffic flows to maintain the end-to-end QoS requirements of diverse applications. Per-flow queuing provides operators a mechanism to manage bandwidth and network resources both at the individual flow level and also as an aggregated resource. Bandwidth transactions can therefore be processed according to network policies and according to the possible impacts they will have on existing service flows. Per-flow queuing also provides the fine-grain monitoring needed so that operators can bill for network resources automatically provisioned by the transactional processes.

Per-flow QoS assigns each packet stream its own queue and provides a guaranteed rate for flows with QoS reservations. It is implemented at the headend using high-performance routers that can classify and treat packet flows in real-time. Operators can track individual flows and ensure that users are within their Service Level Agreements (SLAs) for each session, and they can implement the fine-grain metering required to provide back-end billing systems with metering information on each session.

Since cable networks are based on equipment from multiple vendors, a standards-based approach is required to allow applications to dynamically configure network devices to enable automated provisioning. Operators can treat traffic flows using DOCSIS 1.1 and PacketCable standards on the Hybrid Fiber Coax (HFC) access network and MultiProtocol Label Switching (MPLS) across their own core networks and across the core networks of revenue-sharing partners.

WHY TRANSACTION BASED?

Transaction-based services automate bandwidth management. They eliminate operator intervention and allow transactions to request network resources according to policies defined by the operator or its partners.

Transaction-based services can therefore scale efficiently because they can request additional network resources without operator intervention. Operators do not need to pre-provision every option but can establish network policies that define resource requirements for transactions and trigger the automatic billing for incremental resources.

Effective bandwidth management is the best argument for transaction-based services, since operators do not have to reserve bandwidth that is not in use. Operators can take advantage of statistical multiplexing to more effectively deploy billable bandwidth and support high-value services without creating huge reserves of unused capacity.

The possibilities for transaction-based services are virtually unlimited. Any type of service with QoS requirements is a candidate. This includes streaming audio, streaming video, Video on Demand, telephony, interactive gaming, business applications, video conferencing, and enhanced content delivery.

The core technologies used are DOCSIS 1.1 and PacketCable standards for applying QoS on the access network. PacketCable not only supports telephony but will also support any IP traffic flow with QoS requirements. MPLS is used to provide end-to-end QoS across the core networks of one or more providers, and the ReSource ReserVation Protocol (RSVP) is used to reserve network resources.

IMPACTED NETWORK RESOURCES

Operators must be able to automate the provisioning of resources by autoconfiguring network devices. This includes the cable modem, Multimedia Terminal Adaptor (MTA), or set-top box at the subscriber location.

It also includes the autoprovisioning of the Cable Modem Termination System (CMTS) at the headend. This requires an intelligent, high-performance edge router/CMTS that can deliver QoS on the access network using DOCSIS 1.1 and/or PacketCable standards. The edge router must also be able to deliver end-to-end QoS by serving as an MPLS Label Edge Router (LER) that creates a Label Switched Path (LSP) to the destination. It applies the appropriate MPLS labels that carry the QoS configuration requirements to each intermediate device traveled throughout the metro and core networks.

There are other devices impacted as well. For example, in telephony applications the Call Management Server (CMS) and other Point of Presence (PoP) resources must be configured appropriately to ensure end-to-end QoS.

GAINING A PER-FLOW PERSPECTIVE

The DOCSIS standards are based on the concept of traffic flows and encourage operators to think first about flows and then think about how bandwidth can be adjusted to support flow requirements.

Given the DOCSIS emphasis on flows and diverse application QoS requirements it is critical to look at transactional bandwidth management from a flow perspective, with each data session, application, or voice call treated as an individual traffic flow. Individual flow QoS requirements must then be aggregated to ensure that the sum of their requirements can be accomodated within the available global resource set. Viewing only the aggregated requirements from a network resource allocation perspective can be misleading because this view will not identify potential interactions between individual flows. Thus the transaction processing must consider the requirements for classifying, isolating, policing and enforcing the QoS of the individual traffic flows in addition to looking at the aggregate requirements. In order to make the transaction processing manageable it is essential to provide isolation between the individual flows so that it is simple to calculate the impact of a new flow on existing flows. Per-flow queuing provides an efficient mechanism to implement this isolation.

For applications to dynamically configure bandwidth allocations and device configurations, operators must be able to isolate each flow and have the flexibility throughout the network to treat each flow with the proper QoS parameters. When a telephone goes off hook, an interactive gaming session is launched from the desktop, or a click on a web link launches a business application, a series of steps must be implemented to ensure that the proper QoS treatments are applied on a per-flow basis.

Operators implementing per-flow queuing can offer transactional bandwidth management as elements of new services so that when the subscriber selects the applications the appropriate network resources are allocated automatically. On-demand services can benefit from just-in-time bandwidth provisioning so the customer only pays for resources used. Operators can monitor and meter traffic so that application triggers also feed information into billing and charging systems so that they reap premium payments for premium services. They can also successfully develop relationships with third-party providers of content, applications and services based on fine-grain metering and monitoring to ensure that wholesale partners in turn pay operators for bandwidth used by their subscribers.

Transactional bandwidth management and per-flow QoS therefore open up new opportunities for both retail and wholesale revenue streams. With per-flow QoS, operators can successfully deploy transactional bandwidth services that unleash the broadband potential of HFC infrastructure. They can deploy new pay-for-use services and build closer bonds with subscribers based on increased subscriber abilities to self-select service levels based on their own unique bandwidth requirements.

UNDERSTANDING QOS REQUIREMENTS

QoS control is critical for optimizing the productive use of shared bandwidth on the cable access network. Operators need to be able to allow applications to automatically assign bandwidth appropriately according to service requirements and guaranteed commitment levels. The ability to manage QoS involves four key functions:

- Classification of packets to determine the appropriate service level for each traffic flow
- Policing of traffic to prevent flows from getting higher than agreed upon service levels
- Buffering to ensure that queues are created to contain packets during periods of congestion
- Scheduling to enforce packet handling and actually deliver service end-to-end across access, metropolitan and core networks using Internet standards such as MPLS.

The power and flexibility of policy-based QoS control and measurable QoS levels can support incremental revenue streams from transaction-based services.

BRINGING QOS TO THE ACCESS NETWORK

Network operators are now able to offer transaction-based services via shared cable infrastructure while providing guaranteed QoS levels to each service and user. The DOCSIS 1.1 specifications were developed to define enhancements to the Media Access Control (MAC) protocol of DOCSIS 1.0 to enable more sophisticated access methods over HFC access networks by adding the following:

- Packets are classified into service flows based on their content. Thus each application can be mapped to a unique service flow.
- Network access (upstream and downstream) is scheduled per service flow using one of a number of defined scheduling mechanisms including constant bit rate, real-time polling, non real-time polling and best effort.
- Service flows may be configured through management applications or created and deleted dynamically in response to the starting and stopping of applications.
- Fragmentation of large packets is required to allow low latency services to operate on lower-bandwidth upstream channels.

These features provide the basic tools for transactional bandwidth management. They allow applications to request QoS changes dynamically and allow providers to isolate multiple data streams from each cable modem, set-top box or MTA. DOCSIS 1.1-based systems can therefore potentially deliver the ability to allow dynamic application-specific QoS treatment within the HFC access network for each traffic flow.

END-TO-END FLOW CONTROL

The DOCSIS 1.1 specifications provide QoS for the upstream cable access network. In order for applications to see real benefits, QoS must be provided on an end-to-end basis. Thus the QoS-enabled traffic flows from the access network must be mapped to the QoS mechanism(s) used in the regional or backbone networks.

MPLS can provide the QoS mechanism for the regional network. In an MPLS network a number of paths are established between the end points of the network. Each path can be traffic engineered to provide a defined level of QoS.

The successful combination of these two QoS mechanisms requires that the CMTS must also act as an MPLS edge router to map packets from the DOCSIS flows into the appropriate MPLS paths and vice versa. The mechanisms employed within the CMTS/edge router must maintain the QoS during this transition. The addition of dynamic QoS changes as a result of transactional bandwidth management complicate this problem. Fortunately techniques such as per-flow queuing used in combination with a congestion control scheme such as Longest Queue Pushout Pushout (LQP) can simplify this problem.

Hierarchical Per-Flow Queuing and Longest Queue Pushout for End-to-End QoS

Per-flow queuing assigns each packet stream its own queue and provides a guaranteed service rate for flows with QoS reservations. Those traffic flows that are not assigned prioritization are forwarded in a round robin or fair-share manner. To assign flows without reservations to a queue, a per-flow method known as Stochastic Queuing can be used. The parts of the packet header that are the same for all packets of a flow—such as the source and destination IP addresses and source and destination port numbers—are fed to a hash function that is used to map the packet to a queue. This simplifies the management complexity in a dynamic transaction-based environment because it eliminates the need to pre-configure parameters required by other systems such as the bandwidth shares per-class. Consequently, it also avoids the miss-allocation of resources caused by the varying usage patterns inherent to dynamic bandwidth managed systems and the need to continuously update packet classifiers.

If the system can support more queues than there are flows, than most flows either have their own queues or share them with a small number of other flows. This provides isolation between the different flows in contrast to a class based system where all flows of a particular type (e.g. voice) share common resources.

A suitable congestion control scheme must be selected in addition to the per-flow queuing and scheduling to maintain isolation between flows. LQP is the mechanism that best meets the congestion control requirements in a dynamic system resulting from transactional bandwidth operation. It allocates buffers to the individual flows as required until 100% of the buffer pool is used. When no buffers remain and a new packet is received, LQP discards traffic from the flows which are the longest queues. The scheduling system is transmitting from these per-flow queues at a rate which matches the QoS assigned for each flow. By definition the longest queue is that which is exceeding its allocation by the greatest amount. Therefore, traffic is automatically discarded from those applications which are non-compliant with their SLAs. This occurs without the need to configure congestion control parameters and thus works well in a transaction-based bandwidth management scenario.

The combination of per-flow queuing and LQP congestion control enables a practical system to be built which is capable of operating correctly in a dynamic QoS environment without complex management and monitoring.

The ability to track and schedule based on individual flows can be used to create a system in which scheduling decisions are based on the application flows and the subscribers and service providers with which these flows are associated. Operators can implement transactional bandwidth management services within clearly defined boundaries. With hierarchical per-flow queuing and LQP, operators can schedule packet transmissions and selectively discard packets during congestion based on application needs and based on the SLA conformance of applications, subscribers and service providers.

TREATING FLOWS ACROSS ACCESS, BACKBONE, AND CORE NETWORKS

Hierarchical per-flow queuing allows operators to classify and treat traffic across access, core and metropolitan networks. A next-generation edge router with hierarchical per-flow queuing can inspect multiple fields within packets to determine the appropriate routing and QoS requirements. This requires a powerful QoS routing engine that can inspect packets in real-time and route them across multiple networks according to established network policies.

Transaction-based QoS is invoked by an action by the user. For example, the user could click on a web link or pick up the phone. The cable modem, set-top box, or MTA requests the local resources and QoS-based services using DOCSIS 1.1/PacketCable standards.

A Scenario for Transaction-Based Services

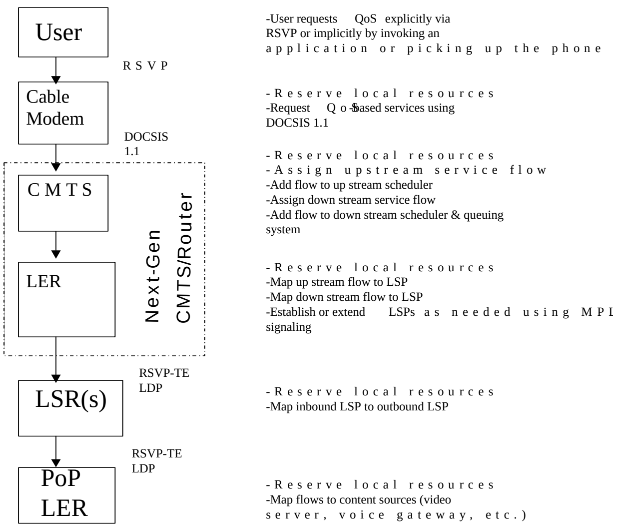

The CMTS portion of an integrated CMTS/edge router at the headend reserves the resources and schedules the traffic flows in the DOCSIS network The LSR portion of the system inspects each flow in real time, maps the packets to an MPLS LSP, applies the appropriate MPLS label and routes the traffic to an MPLS Label Switched Router (LSR) in the core network. The LSRs then switch the packets across the MPLS network to the corresponding LERs at the points of presence for the content sources.

If a flow requires a new LSP (or changes to an existing LSP) then RSVP Traffic Engineering (TE) extensions or Label Distribution Protocol (LDP) messages are used to signal the setup of the MPLS path.

Authorization and Billing

Allowing multiple service providers to deliver transaction-based services over a shared access network requires extensive features for authorization, reconciliation and billing.

Operators need the flexibility to develop policies that determine whether a transaction is allowed at each point in the path and to implement admission control policies that determine which applications are prioritized during times of congestion.

They also need to determine what data to collect, where to collect it, and how to implement reconciliation between service providers. Operators need to maintain detailed accounting information on QoS usage to ensure that application commitments are enforced.

Per-flow queuing allows operators to observe and manage individual IP service flows. Each session can be carefully tracked, and per-flow metering information can be automatically exported to third-party account or mediation applications using the IP Detail Record (IPDR) format or Call Detail Record (CDR) formats.

Granular observability into traffic allows operators to implement enhanced billing applications that reach far beyond traditional flat-rate billing to allow operators to deliver—and accurately bill for—premium dynamic services that automatically increase bandwidth flows in response to application requirements. They can therefore deliver highly granular bandwidth management capabilities with detailed reporting and accounting.

TRANSACTIONAL BANDWIDTH MANAGEMENT WITH PER-FLOW QOS

Operators can now implement transactional bandwidth with per-flow QoS by deploying intelligent edge routers that can classify and treat traffic flows to ensure end-to-end QoS across cable access, metro, and core networks.

A standards-based approach allows operators to build infrastructure that can support high-value, transaction-based services that help them increase revenues, market share, and profits.

###

About the Author

Gerry White serves as Chief Technologist for the Network Infrastructure Group of Motorola Broadband Communications Sector, and he was previously CTO of RiverDelta Networks, Arris Interactive and LanCity. He has presented papers and spoken at dozens of cable industry events worldwide and has published many trade press articles on cable industry issues. White is the co-author of several patents and articles on data communications technology, and he holds a BSc. (Honors) from University College in London.

Author Contact Information

Gerry White
Chief Technologist
Network Infrastructure Group
Motorola Broadband Communications Sector
3 Highwood Drive
Tewksbury, MA 01876
(978) 858-2300
gwhite@motorola.com