PQoS – Parameterized Quality of Service

White Paper

Abstract

The essential promise of MoCA—no new wires, no service calls and no interference with other networks or consumer electronic devices—remains intact with each upgrade to the standard.

MoCA 1.1, in addition to packet aggregation that yields net throughputs of 175 Mbps, and 16-node availability, also includes Parameterized Quality of service (PQoS). PQoS ensures that asynchronous data that need only be delivered on a best effort basis, for instance, does not interfere with the delivery of time sensitive multimedia data streams. This is a critical challenge for system operators as they address home entertainment networking usage models evolving from simple transfer and download of best effort data (Internet access for example) to a media centric model where data is streamed around the home and must be delivered in real-time.

The MoCA 1.1 standard provides a parameterized QoS mechanism that enables operators to either leverage the efforts of industry standard high-level PQoS frameworks or create a custom QoS solution using custom designed higher layer QoS services.

MoCA is the first home entertainment network standard to embed PQoS.

Quality of Service Challenges in Home Entertainment Networks

Sharing Internet access and transferring files between personal computers is a subscriber's traditional notion of the home network. However, as more multimedia-based applications become ever common, these networks must support a deeper and more complicated usage model that requires real-time streaming of high definition (HD) video and entertainment. Consumers now expect to be able to watch digital video recorded (DVR) TV shows, access video on demand (VOD) services, stream music from their personal computers and play interactive games on the Internet, in real-time and from virtually any room in the home.

System operators face the challenge of supporting these real-time, multimedia streaming applications simultaneous with Internet access traffic while maintaining coexistence with already existing services such as analog/digital/IPTV and VoD, within the same network environment. Critical to effectively serving all these data flows is a method to ensure that each application is guaranteed the bandwidth and minimal latency necessary to provide a satisfactory user experience.

Figure 1 illustrates how a home entertainment network can support a variety of devices, each with its own requirements and performance expectations, over coaxial cable. The network could be carrying video, data and voice, and consist of a mixture of QoS- and non-QoS- enabled devices. A network this complex poses significant network bandwidth management challenges especially when multiple real-time streams available within the LAN or the WAN need to be transmitted to devices with varying capabilities.

An unsatisfactory user experience can be costly to system operators. During data transmissions that are delivered on a best effort basis such as Internet page downloads and file transfers, subscribers are accustomed to delays that result in a relatively small impact to their experience. But delays or latency in a video transmission is unacceptable and adversely affect the entire experience. A disruption to the flow of streaming video or audio can result in stuttering playback, blocky video frames and a complete loss of audio, prompting a service call from the subscriber. In order to avoid these situations and enhance the viewers video and entertainment experience, quality of service (QoS) methods are used to manage the priorities of the different traffic flows on the network.

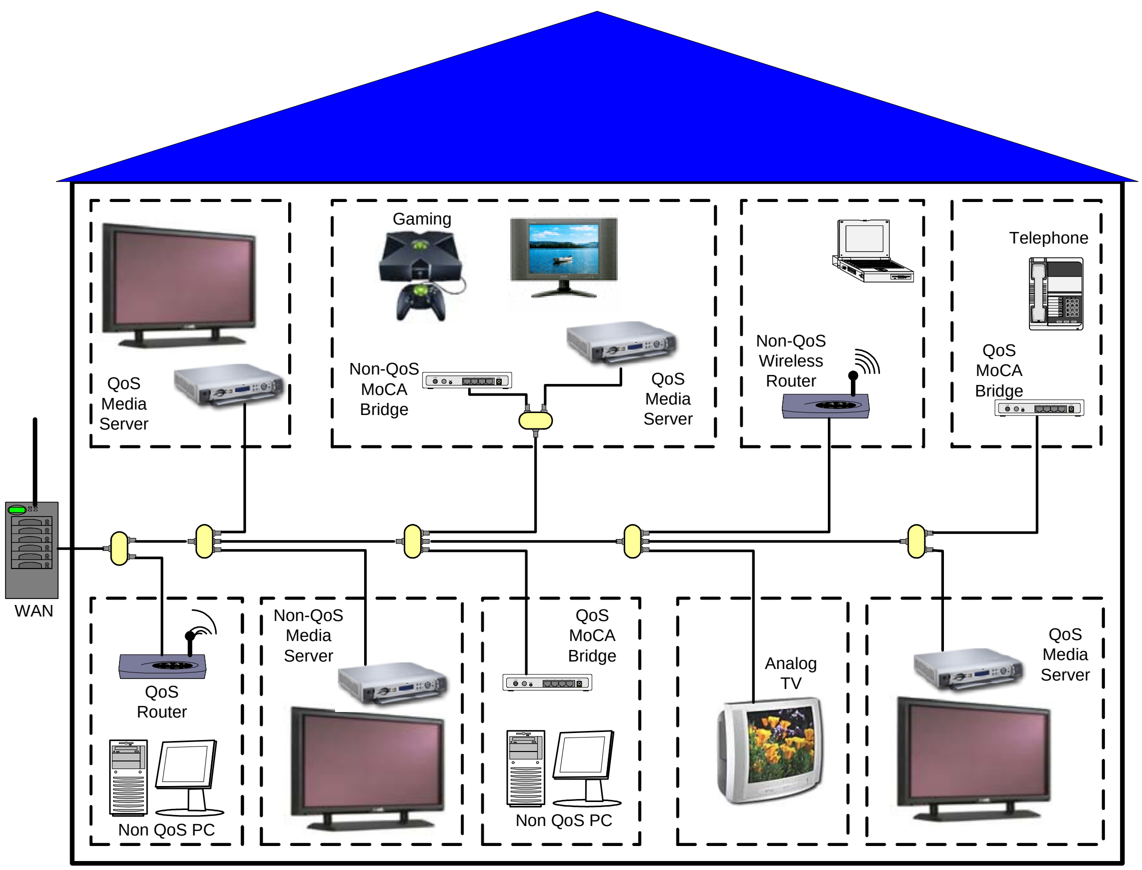

Figure 1: Home Entertainment Network

Today, 802.1p provides a method to prioritize traffic on the network. Unfortunately, this mechanism does not make distinctions between types of data nor does it guarantee bandwidth. It only provides the network with a method to weigh the importance of one priority against that of another priority.

Asynchronous internet data tagged with a priority of 5 would be treated the same as video data tagged with a priority of 5. The result could still lead to an operator's high quality video service being impacted by regular data.

Even multiple video streams tagged as high priority are not immune from interfering with each other in an 802.1p QoS scheme, as routers simply queue up all of these packets in its high priority buffer to be sent as soon as possible. But if the bandwidth is insufficient to handle a large amount of high priority traffic, packets will be dropped. As end consumers add third-party products and applications available through retail channels, operators lose more control of their video network and bandwidth conflicts between video and data intensify.

Another potential solution is the time honored tradition of throwing bandwidth at the QoS problem, also known as the brute force approach. However, this is at best a stopgap measure as an applications thirst for higher data rates continue to grow ever larger, eventually using up any additional bandwidth and still leaving the problem of not being able to use the available bandwidth effectively.

In addition, adding more bandwidth doesn't take into account that simple network file transfers can consume all of the bandwidth in transient bursts. Operators and subscribers require a QoS mechanism that protects their premium content yet is flexible enough to work in a mixed mode environment.

Instituting a parameterized quality of service (PQoS) mechanism avoids costly operator support and unsatisfactory user experiences. PQoS provides true guaranteed performance by reserving the required network bandwidth based on the characteristics or parameters of that data requesting access to the network. The ability to limit over subscription of the network is enhanced through an intelligent negotiation between devices which determines whether the network can accept or reject the requested bandwidth.

PQoS mitigates bandwidth conflicts and prioritization as it ensures there is enough bandwidth within the home network to accommodate worst case loading scenarios. A minimum of 100 Mbps (and more preferably) net throughput, statistically and independently validated, is required for immunity to dramatic bandwidth interference and is more than adequate to handle multiple HD streams within an entertainment network. If heavy loads of asynchronous data traffic are introduced into the network, PQoS guarantees that video flows have priority over flows that do not require real-time streaming.

Economic benefits abound as service call and truck roll costs, which can run north of $250, are reduced if not eliminated by providing a robust network that can accommodate time sensitive data streams simultaneous with large ftp sessions, for example, without negatively impacting the user's experience.

A significant benefit that PQoS provides to equipment manufacturers is lower costs. With guarantees of network bandwidth and latency, client receivers can be cost reduced by eliminating the need for hard disk drives and large amounts of buffer memory that were once required to work around network uncertainties.

Upper Layer QoS Framework

In order to offer comprehensive support for all types of data services and a wide variety of end devices, MoCA has adopted a QoS framework intended to easily integrate with standardized higher layer QoS services. These upper layer QoS services are responsible for the creation and tear down of PQoS flows.

By leveraging a standardized QoS framework, layer 3 and above compatible applications can enable PQoS MoCA networks using existing QoS management services. However, MoCA contains a simple set of low-level services that provide the flexibility to also support proprietary implementations of vendor/operator QoS management architectures.

The type of QoS method used for each segment of the network is based on the capabilities of the devices attached to each end. If a QoS aware application requests a prioritized QoS traffic flow, it is the responsibility of the QoS management service to attempt setting up a prioritized traffic stream through the segments. Similarly, if an application requests PQoS, the QoS management service will attempt to setup a parameterized traffic stream. In the case where hybrid QoS is requested, the manager will attempt to establish prioritized QoS on segments that support priority based traffic flows and establish PQoS on segments that support parameterized traffic flows.

How It Works

PQoS allows systems operators to assign higher priority or walled garden to streaming video flows above any other traffic through higher level device negotiation. Also sometimes referred to as reserved bandwidth, PQoS allows the entertainment network to precisely define its quality of service needs, reserving just enough node-level and network-level resources to guarantee the necessary bandwidth.

A robust layer 2 device will automatically implement a set of PQoS autonomous behaviors. These low-level autonomous actions are used to build PQoS capability at higher network levels. Three interfaces are provided and each is responsible for the following:

- QoS Operation Interface
    - Create uncast or broadcast flows
    - Modify flow attributes for existing flows without disrupting the stream
    - List flows on each node for which that node is ingress and query their attributes
    - Delete any flows
- QoS Event Detection Interface
    - Detect whenever a node is added or removed from the network
    - Detect when an ingress flow is added, deleted or updated
    - Detect when parameterized bandwidth has been exceeded
    - Detect when the network has enabled or disabled QoS
- General MoCA Interface
    - Provides APIs/services to query about network topology information such as number of nodes, their capabilities and GUIDs

Overall bandwidth can be thought of as being divided into two parts, parameterized flows and asynchronous flows. The maximum bandwidth allocated for parameterized flows is typically set to a certain percentage of the overall bandwidth. The remaining bandwidth is used for all asynchronous traffic, including prioritized traffic, best effort traffic and link maintenance.

Higher level applications create a flow by making a request to the PQoS interface which includes the peak bandwidth, nominal packet size and other parameters. These parameters are then converted into the scheduling elements necessary to calculate if the bandwidth is available. If the bandwidth is available, the flow is authorized and the bandwidth is taken out of the general resource pool and allocated to that specific flow.

A flow is identified and assigned by a higher network layer by using the destination address field which is set to be a multicast MAC address and the VLAN tags corresponding to user priority of 4 or 5. Flows are terminated and the resources given back to the network either when the user terminates the session or an optional lease timer causes an automatic expiration of the flow. In the event the network bandwidth requested based on the parameters provided is not available, the flow request is denied and data traffic will not be permitted between those nodes using the PQoS mechanism. In the process of denying the requested flow the amount of bandwidth which is available, based on the previous parameters, is also passed to the requestor. This enables a node to still create a flow if it can abide by the bandwidth limits.

PQoS Is In The House!

PQoS enables the home owner to effectively mix and match different types of traffic without sacrificing performance. It is flexible enough for system operators to implement a either a standardized PQoS model or a proprietary scheme embedded within MoCA 1.1.

Enabling higher performance and the ability to segment traffic yields great benefits to both the service provider and the end consumer.