AN EXTENSIBLE QOS ARCHITECTURE FOR A HETEROGENEOUS NETWORK INFRASTRUCTURE TO SUPPORT BUSINESS SERVICES

Srividya Iyer, Dr. Nagesh Nandiraju, Dr. Sebnem Zorlu-Ozer
Motorola – Access Networks

Abstract

Cable operators have been deploying business services for some time now. The nature of these services demand higher bandwidth, better management and complex service level guarantees. In order to meet the demands of the customer, the heterogeneous networks (DOCSIS, Ethernet, PON, Wireless) deployed by the cable operators have to be integrated to work seamlessly to support these services.

The technologies in use for delivering business services are predominately packet based and hence require different Quality of Service (QoS) mechanisms to ensure proper delivery of services. However, the QoS mechanisms defined for these technologies have evolved independently and interoperation in a multi-technology network environment can be difficult. Often the network architects and engineers managing these networks face tremendous challenges in translating the QoS definitions and rules from one network to another, thus making it difficult to provide a seamless QoS experience for the users.

In this paper, we propose an extensible QoS architecture that will support the heterogeneous network infrastructure. Specifically, we will address bandwidth reservation policies, priority queuing and the relationship between the Layer 2 and Layer 3 (IP) QoS mechanisms in a multi-technology network.

INTRODUCTION

Business Services is rapidly becoming the main growth area for cable operators. As guaranteed services have become commonplace for this customer base, it has become necessary for the operators to address the customer's quality of service concerns.
Quality of service is defined as the collective effect of service performances, which determine the degree of satisfaction of a user of the service (ITU-T Rec. E.800).
Quality of Service in general is determined by the network performance for a given service, measured by throughput, delay, jitter and packet loss. Translating a customer requirement into a Network QoS to satisfy the customer's perception of Quality is illustrated in the following diagram [1]

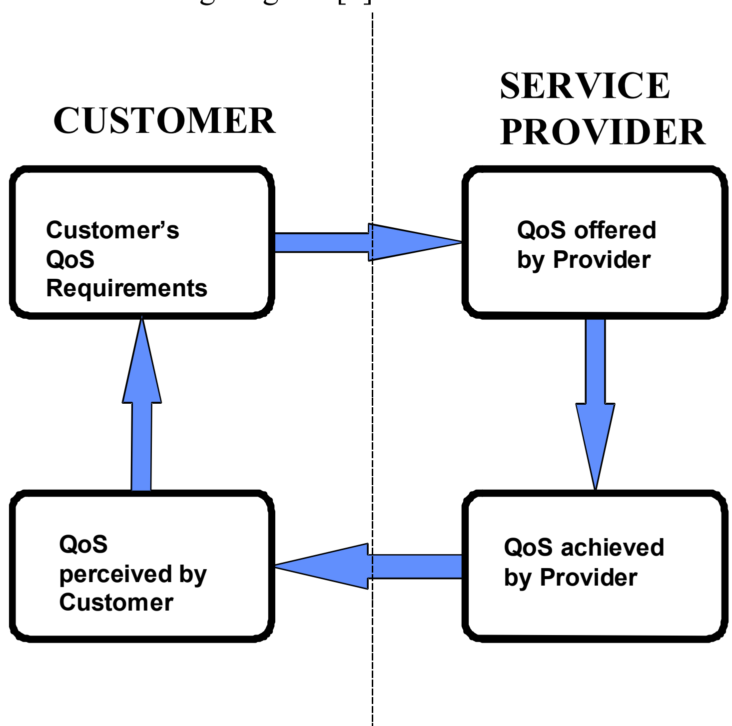

Figure 1-The four viewpoints of QoS

Providing a high level of QoS is possible if the networks are designed uniformly and are under a single management entity. In reality, the cable operator's networks are usually designed taking into account the services required, cost and distance making the networks heterogeneous in nature. Heterogeneity exists in technology, transmission media, applications, user devices and individual vendor implementations. This heterogeneity makes providing end to end QoS a challenge due to different QoS implementations and lack of communication across these network domains.

The rest of the paper is organized as follows: Section 2 provides a brief overview of the various QoS mechanisms and a mapping between different types of applications.
Section 3 describes the challenges with providing QoS in heterogeneous network architecture. Section 4 describes the solution proposed for providing a consistent QoS across multiple network domains. Section 5 concludes the paper.

SECTION 2 - KEY QOS CONCEPTS

Business services typically consist of varied applications with diverse requirements. Each type of service has a different set of bandwidth requirements and tolerance for latency, jitter and packet loss. The end user requirements, usually defined in subjective terms needs to be translated into Network performance parameters. In order to achieve the necessary performance parameters, QoS mechanisms in packet based networks are implemented either at Layer 2 (MAC) or at Layer 3 (Network or IP). The QoS mechanisms at Layer 2 operate in the individual domains (Ethernet, Wi-Fi, DOCSIS, PON), while the Layer 3 QoS mechanisms provide End to End QoS within the IP domain.

In order to achieve the desired Quality of Service characteristics, some of the QoS mechanisms implemented include Packet classification, Queuing, Bandwidth reservation and Traffic conditioning.

Packet classification – Provides the capability to classify the network traffic into multiple priority levels or classes of service. Packet Classification can be done either at the Layer 2 or Layer 3. Typically a “classifier” is used to classify packets.

Priority Queuing – Usually implemented to avoid congestion in the network by placing the packets in buffers till bandwidth becomes available. Priority queuing uses the various packet classifications and, based on their priority, places them in various queues (high, medium, low etc.) Examples of priority queuing include 802.1p queues and the Diffserv forwarding classes.

Bandwidth Reservation – Usually employed when a network is transporting traffic that requires minimum bandwidth guarantees. In such scenarios, each application will receive a predetermined minimum and maximum bandwidth allocation. The service provider needs to provide the customer the best possible QoS while conserving the network resources. This requires some kind of traffic estimation to prevent either under utilizing the bandwidth or over subscription.

Traffic conditioning – The process of metering, marking, shaping or dropping traffic. Traffic conditioning provides the enforcement of the traffic profiles defined for each of the services.

QoS mechanisms at Layer 2 vary according to the technology implemented. The four commonly used technologies by the cable operators include Ethernet (bridged), DOCSIS, Wi-Fi and PON.

Ethernet uses the 802.1p priority bits to classify the traffic into various classes. The priority queuing uses these traffic classifications to place different classes of service into different queues.

DOCSIS uses a concept of service flows to classify packets into a unidirectional flow which is then shaped, policed and prioritized according to the QoS parameter definitions. Multiple service flows can exist for each Cable Modem and grouping of service flow properties is also facilitated through the definition of a service class.

The QoS mechanisms defined for 802.11 networks include an additional coordination function called HCF (Hybrid Coordination Function) on top of DCF (Distributed Coordination Function) and PCF (Point Coordination Function). The HCF enhances QoS provisioning during both contention and contention free periods by using EDCA (Enhanced Distributed Channel Access) and HCCA (HCF Controlled Channel Access) respectively. Due to higher complexity and complications caused by overlapping stations and unlicensed spectrum use, HCCA is not an industry choice. The Wi-Fi alliance WMM program is based on the EDCA method and was initiated as the market needed a QoS solution before 802.11e was ratified. WMM defines four access categories (voice, video, best effort, and background) that are used to prioritize traffic to provide enhanced multimedia support. EDCA defines four AC (Access Categories) to differentiate different services. This requires mapping from UP (User Priority) to AC. 802.11e recommends UP to AC mapping, however current approved standards do not provide a uniform implementation for vertical (between Layer 2 and 3) and horizontal (between various network domains at Layer 2) QoS mappings. The ongoing 802.11u standardization effort aims to standardize the "information transfer from external networks using QoS mapping".

In Passive Optical Networks (PONs), Gigabit-PON (ITU standard G.984) introduces the concept of Traffic Containers (TCONTs) to classify the traffic and service them according to the QoS definitions. In Ethernet PON (IEEE 802.3ah/av) multiple Link Layer Identifier(LLID) can be used to classify the traffic into different classes and service them according to the negotiated QoS definitions.

QoS architecture frameworks at Layer 3 consists of either a differentiated services model (DiffServ IETF RFC 2474/2475), where the traffic is classified and the frames are treated with different priority based on information carried in the frame header, or a reservation model (IntServ IETF RFC 1633), where a signaling is used per session to reserve resources.

SECTION 3 - CHALENGES OF IMPLEMENTING QOS IN A HETEROGENEOUS NETWORK

There are several challenges facing the service provider in providing a high level of QoS across heterogeneous network architecture.

First and foremost, enabling QoS has been focused on mechanisms and protocols in individual network domains (eg wireless or cable access) or even individual network elements. This provides a high level of QoS within that domain or network element, but lack of communication between the various domains makes it harder for the QoS implementations to cross boundaries. That leaves the network architects and engineers managing the networks with the task of implementing the disparate QoS policies within each of the devices in each of the networks. In addition, the devices provide a multitude of options to the network architect. These options might not be the same across networks or even across different vendor devices within a network. This makes the QoS implementation error prone and inconsistent.

Almost all the QoS mechanisms that exist are defined statically to enable QoS and do not automatically adjust to the traffic conditions that exist in the network. To solve this issue, the network architects routinely over provision the network leading to wasted bandwidth. Conversely, where bandwidth is at a premium, the network is oversubscribed leading to lower priority services being denied access. Also, it is to be noted that over provisioning the network doesn't always guarantee required QoS to the end user as evidenced by the peer to peer applications consuming a large portion of the bandwidth and causing degradation of other services.

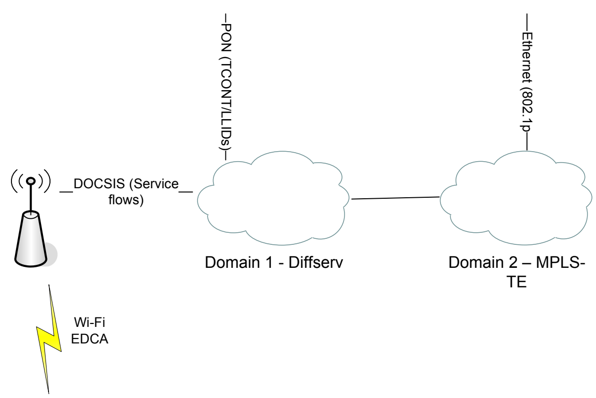

Figure 2 - An example of QoS mechanisms in a Heterogeneous Network

Several standard bodies have defined service classes and their packet classification schemes. However, mapping these service classes to the individual QoS mechanisms in each of the network domains still remains an issue for most service providers. Since there is no fixed number of service classes supported in each domain (Ethernet 802.1p supports up to 8 service classes and Wi-Fi APs support four Access Classes) and no enforced mapping between the service classes and the QoS mechanisms, it is up to the person defining these QoS mappings to select the best possible one, thus leading to inconsistent behavior across the network.

Typically QoS can be provided at the Layer 2 or the Layer 3 level. The current implementations usually mix Layer 2 and Layer 3 QoS mechanisms (For example, a device can implement the 802.1p CoS with a Diffserv AF PHB). The traffic class markings between packets can vary depending on the originating or forwarding device and can be at Layer 2 or 3 or both. The device would then use the Layer 2or Layer 3 class markings depending on which one is available or based on network configuration (trust Layer 2 or 3) . The challenge in this case is to provide a consistent mapping between the Layer 2 and Layer 3 mechanisms. As indicated earlier, they evolved independently and thus there is no standard way of mapping the Layer 2 to the Layer 3 QoS, leaving it to individual implementations.

SECTION 4 - AN EXTENSIBLE QOS IMPLEMENTATION

Several consortiums and research bodies are in the process of defining end to end QoS architectures [2]. These architectures define a method to provide end to end QoS and might benefit certain new network builds. They however require an exhaustive rework and re-architecture to function within a large existing network infrastructure.

Implementing a network-wide end to end QoS architecture is time consuming, requires a lot of resources and is disruptive to the existing operations. In order to address the challenges outlined in the previous section, we are proposing an extensible QoS architecture that would use the existing QoS mechanisms in each of the domains optimally and with minimum manual configuration, while providing a seamless QoS experience to the user.

Service classification and Traffic category mapping

In any network that provides more than a Best Effort service, the first step for the network architect is to classify the types of services that will be provided over the network. Having a very granular service classification provides the best possible QoS for each class, however in a multi-technology network, this becomes cumbersome and scalability becomes an issue.

The first step is to classify the services into broad categories that share similar requirements. IETF and 3GPP both define four service classes and can be used as a reference, or the service providers can define their own set of classes. The key is to define service classes and the corresponding traffic category mapping that enable a consistent behavior across multiple domains. The following table illustrates a subset of service classes and the recommended (or commonly implemented) Layer 2/Layer 3 QoS schemes available across a sampling of the different network architectures.

Traffic Category	**Ethernet, Wi-Fi DOCSIS,PON**	**Diffserv Layer 3**
Real Time Traffic (Voice)	P bit 6, 7 AC3 UGS,TCONT1	EF
Interactive (Video)	P bit 5 AC2 RTPS,TCONT2	AF
High Availability (Data)	P bit 3, 4 AC1 nRTPS,TCONT3	AF
Best Effort (Data)	P bit 0,1, 2 AC0 BE,TCONT4	Best Effort

Table 1 – Traffic category/QoS mapping

Layer 2 to Layer 3 mapping

With the advent of heterogeneous networks and Layer 2 Services (in order to avoid expensive Layer 3 equipment), several Layer 2 networks will be interconnected between the Layer 3 core networks. This introduces interoperability issues between the Layer 2 and Layer 3 QoS mechanisms.

Although the IETF states that the Diffserv model can be extended to support Layer 2 architecture, there is no defined behavior or mapping between the Layer 2 and Layer 3 QoS mechanisms. In our model, we propose implementing a consistent set of traffic category mapping with the Layer 2 QoS mechanisms available and updating this information across all the network elements. This is in order to avoid inconsistencies that arise from the manual configuration of QoS profiles in each network element. A specific implementation of this mapping is hard to define as there are variations within the network domains and network elements.

Queuing

One of the key service differentiators used in supporting multiple classes of service is priority queuing. The priority queues will usually be implemented separately for both ingress and egress traffic. However priority queuing does not provide bandwidth guarantees and deterministic latency, jitter and packet loss characteristics. In addition, each network element supports two or more types of priority queuing algorithms and different numbers of priority queues thus causing inconsistent treatment of service classes across various network elements. That being said, priority queuing is widely used and an integral part of providing service differentiation. Two key elements in using the priority queuing efficiently are mapping the service classes to the right queue and selecting the right algorithm. Neither the Strict Priority (SP) nor the Weighted Fair Queuing (WFQ) /Weighted Round Robin (WRR) provide the fairness required for all traffic conditions. In the Strict Priority queue, the highest priority traffic always gets through and, depending on the traffic load, the lower priority queues can be starved for bandwidth. In the WFQ/WRR scenario, the weights that are assigned are static and can either not provide the required QoS or waste bandwidth depending on how many different service classes are accessing the network at a given time.

In our QoS architecture, we propose a scheme where the type of queuing employed at the ingress and egress and the weight assigned to each of the queues is based on the traffic received on each of the queues over a period of time. The time interval for sampling of the traffic can be predetermined by the network element or configured. This eliminates the problem where the high priority traffic always gains unfair advantage, while the low priority queues are denied access. This also eliminates the wasted bandwidth with statically configured queues. In addition, the minimum and maximum rate of each queue will be adjusted dynamically based on the traffic conditions.

<u>Bandwidth reservation</u>

In most services that require high Quality of Service, bandwidth reservation is one of the key factors affecting the QoS and additionally has an effect on the delay, jitter and packet loss characteristics.

IETF proposed IntServ with RSVP to reserve bandwidth for each flow end to end. This approach, while providing the optimal bandwidth for each flow, is not scalable and hence not supported in most network elements. On the other hand, due to its simplicity and scalability, DiffServ is more widely supported. Diffserv only provides bandwidth guarantees for aggregate flows and defines a set of Per Hop Behaviors (PHBs).

Most of the network elements support some level of DiffServ functionality. For example, A DiffServ compliant (DS) node that supports the four Assured Forwarding (AF) traffic classes must allocate a configurable minimum amount of forwarding resources (buffers and bandwidth) to each AF class. The minimum bandwidth is pre-allocated for each AF class and is equal to its corresponding Committed Information Rate (CIR). This pre-allocation of bandwidth has to be done on all the nodes supporting the AF classes. In addition, the bandwidth availability and traffic patterns in each of the node vary. Pre-allocating the CIR bandwidth for each AF class might be too much or too little.
In order to overcome the challenge of pre-allocating bandwidth at every node for each of the AF classes, we propose a dynamic bandwidth allocation scheme at each node based on the service classes defined. This is done as a two step process. Each node maintains a count of aggregate traffic received for each AF class and the Best Effort services. In addition, each node supporting a DS function will query the adjacent node for bandwidth usage to support dynamic bandwidth allocation and avoid congestion.

<u>Communication between multiple network domains</u>

With the evolution of heterogeneous networks, there is a need for multiple network domains to interoperate and provide seamless

QoS experience to the end user. Two of the main issues with interoperability of QoS between these domains is the lack of consistent QoS implementations and communication between different entities. The QoS implementations have evolved independently and it would not be practical to change the existing mechanisms. Providing end to end signaling across these networks is also not feasible due to scalability issues. In order for services traversing these domains to achieve high level of QoS, we propose a communication path between the nodes residing at the edge of each domain. These edge nodes would ensure that the QoS required for a service traversing that domain can be satisfied by the resources available.

SECTION 5 - CONCLUSIONS

In this paper, we propose a QoS implementation scheme that utilizes the existing Layer 2 (MAC) and Layer 3 (IP) QoS mechanisms, while providing a consistent QoS to the end user across multiple network domains. This approach also minimizes the number of static QoS profile configurations that are needed in each network element. The key to achieving QoS interoperability across multiple network domains is to recognize and reconcile the different QoS implementations that exist today. There is also a need to focus on the end user service requirements while optimizing the network resource utilization.

REFERENCES

[1] ITU G.1000 Standard, Communications quality of service: A framework and definitions, 2001.

[2] B. Iancu, V. Dadarlat, A. Peculea, “End-to-End QoS Frameworks for Heterogeneous Networks - A Survey”.