

MRV Communications

OS9000 - Traffic Management QoS Services




Table of Contents

INTRODUCTION	3
QUALITY OF SERVICE (QoS) NETWORK CHALLENGES	3
<i>Applications Classes of Service</i>	4
OS9000 QoS ARCHITECTURE	5
<i>Classification of Traffic – Defining a Flow and Action Rules</i>	6
QoS AT ETHERNET, IP AND MPLS LAYERS	8
TRAFFIC MANAGEMENT SERVICES	8
ENHANCED HARD-QoS MECHANISM – MPLS CONTROL PLANE	9
CONCLUSION	9

Introduction

The OptiSwitch 9000 offers enhanced optical Ethernet switched services for both business and the residential markets with triple-play (voice, video and data) networks. Nowadays, a provider access network is de-facto based on Ethernet devices with several layers of aggregation. This concept brings about a distributed model of Ethernet switched services and correlating requirements for different types of aggregation devices, with small to medium densities that can fit the first aggregation, the one closest to services and users. Broadband services must be delivered with appropriate level of QoS throughout the network, while efficiently optimizing the bandwidth use on the first mile and second mile aggregation to reduce operating costs. The objective of this paper is to outline applications requirements and the related network QoS implementations provisioned by the OS9000 series, at the first and second mile aggregation points of a provider's network.



Quality of Service (QoS) Network Challenges

Today and in the foreseeable future, IP application services combined with Ethernet transport are the ultimate choice. However, unaided IP Ethernet networks only provide "best-effort" QoS.

Since IP and Ethernet are connectionless technologies, traffic can be lost, dropped, or never delivered for several reasons - especially when the network gets busy. This implies that in order to serve properly QoS-oriented applications, traffic should be marked, classified, scheduled and policed based on committed rates.

From the overall network service perspective, QoS provides end-to-end traffic control over point-to-point and multipoint connectivity services, so that users' applications like VoIP and streaming video can be properly served according to the allowable quality requirements such as committed bandwidth, latency, jitter and packet loss.

Traffic Type	Transport	Bandwidth	Latency	Jitter	Packet Loss	Packet Sizes	Priority
Data	TCP	Bursty	Resilient to	Resilient to	Resilient to	Variable	Low
Voice	UDP	Low	Low	Lowest	Medium	Small	High
Video	UDP	Highest	Low	Low	Lowest	Large	High

To comply with the service quality requirements, traffic that originated from an enterprise LAN or a residential customer, at the application layer (typically IP) should coordinate QoS traffic control with carrier's transport level QoS.

Differentiated service delivery with respective service levels starts at a basic point of the customer's demarcation, and continues to the interior aggregation and core network. It is absolutely critical to define how the QoS mechanisms can be applied at various points in the network to achieve appropriate performance characteristics.

By assigning different service levels to traffic flows, applications can be granted the bandwidth and network priority required in synchronization with all network element.

Applications Classes of Service

The combination of IP and Ethernet in network architectures provides traffic class and QoS awareness using the Ethernet IEEE802.1p, IP DiffServ Code Points (DSCP), or MPLS EXP bits (E/L-LSP) and other indicators to differentiate between classes. Accordingly, it is possible to treat some classes of traffic differently to ensure they match their QoS requirements. There exist up to 8 defined service levels (classes) that network providers can map to correlating applications. These classes have a direct correlation to the 8 levels of priority that originate from the number of bits that can be marked.

This implies that when certain traffic flows are marked with high priority and others with lower priority, if congestion occurs, the lower priority traffic will be dropped first or remarked based on a policy that will be defined.

Service providers generally define the Class of Service as Service Levels and these are one of the considerations when establishing contractual agreements to be offered to customers. Each service level can process mixed traffic with different QoS and traffic profiles (delay, jitter, loss, bandwidth and availability).

Due to the wide range of applications in today's broadband networks, the following three different service grades are the most common:

- **Real time** - Most rigid requirements demanded from the network: This category includes management communications between network elements (typically marked with value 7), and real-time applications like VoIP, Video conferencing and other delay sensitive applications (typically marked with values 5 or 6).
- **High priority & Streaming** - A nearly real time transport of voice and video or business ERP applications.
- **Best Effort** - like e-mail and web browsing that require only the network's best effort, i.e. the lowest priority.

Real-time

VoIP and Video-conferencing are defined as real-time applications. For the real-time class, the requirements for minimal delay and delay jitter are very strict; delay and appropriate bandwidth are absolutely critical: with too high a delay, interactive communication will not be possible. Real-time traffic is transmitted via UDP. Because of the real-time character, a retransmission of lost packets is not useful, because packet retransmission will only increase delay tremendously. Since real-time traffic is very delay sensitive, a Strict Priority (SP) service will be the most desirable. Thus, guaranteed minimal delay and delivery time are ensured.

High-priority Streaming

High priority refers to business applications like ERP, CRM and others. Video and Audio streaming include Video-on-Demand and video broadcast applications known as IP-TV. In such a class, jitter is compensated by application play-out buffer with an impact on the delay. Some streaming applications also detect the network status and buffer an appropriate amount of data. Streaming is particularly oriented for audio and video transmission even for stored data as live feed. The streaming information is normally transported via UDP.

Best Effort

For non real-time applications there is no determined requirement in delay. Non real-time applications examples are e-mail, FTP, HTTP etc. Non real-time traffic is transported with TCP. TCP is a reliable protocol, which transmits each byte of information.

This can take a longer time when retransmissions occur. In today's IP networks, these are the standard applications and they require only best effort.


To ensure abovementioned services, MRV offers only standard non-proprietary solutions when designing and implementing Metro Ethernet networks. Such strategy simplifies management and interoperability within multi-vendor Metro networks with a variety of network elements and a wide array of services and applications.

OS9000 QoS Architecture

The OptiSwitch® 9000 series offers intelligent QoS services at all its physical layer ports without performance degradation at line rate. The OS9000 enforces QoS parameters on a per-service basis and can be positioned at the access and aggregation points of the first and second mile of a provider's network. The network building blocks start with the customer's demarcation, known as the User Network Interface (UNI), that is connected by the access device and aggregated by the interior provider network, called first and second mile aggregation network. In a DiffServ model, the UNI acts as a DiffServ access inbound point, and the aggregation of the network acts as interior Per Hop Behavior DiffServ aggregation. Each network element differentiates the various application flows by classifying the flows at the access (customer's demarcation) with QoS marking. The interior (aggregation) network provides queuing and scheduling for the aggregate flows.



Classified flows can be defined on traversed services/flows, whether it is Voice over IP, video over IP or just best effort data, and are capable of matching each service criteria. Each service is isolated and mapped with the relevant QoS profile, including correlating bandwidth control and service level. Today's common techniques for Metro Ethernet services have similarities with legacy Frame Relay. Customers or applications are policed and shaped based on subscriber's VLAN or Layer 2-3-4 packet headers. Such bandwidth service is offered on each of the traversing flows in the OS9000 together with Committed Information Rate (CIR) and Excesses/Peak Information Rate (EIR) with a flexible granularity of 64 Kbps. The OS9000 offers significant enhancement for traffic management by offering dynamic bandwidth sharing between subscribers. This feature enables the preservation of a SLA for active subscribers/applications like guaranteed bandwidth, and sharing of the unused bandwidth from non-active subscribers/application.



The OS9000 QoS architecture leverages a variety of mechanisms to ensure end-to-end service, based on the following functional models:


1. Virtual LAN (VLAN) - tagging, double tagging and MPLS label mapping
2. Class of Service - classification of traffic to micro flows based on Layer 2-3-4 headers
3. Marking/remarking Ethernet headers - IEEE802.1p
4. Marking/remarking IP ToS headers - DSCP
5. DiffServ access and interior mechanisms - Assured and Expediting forwarding
6. Scheduling - queuing & buffering for appropriate QoS treatment
7. Policing - single/dual rate 3 colors per logical flow and physical interface
8. Egress shaping - per physical interface and individual queue
9. Congestion avoidance - layer 4 flow control RED/WRED within each queue

Classification of Traffic - Defining a Flow and Action Rules

The common method for identifying flows is the Access Control List (ACL) mechanism.

The OS9000 is capable of recognizing a series of frames as belonging to a specific flow and of applying policy-based traffic classification and prioritization based on the following parameters:

- ▶ VLAN
- ▶ Protocol (ICMP, IGMP etc.)
- ▶ IP source / destination
- ▶ TCP/UDP source / destination
- ▶ IP ToS
- ▶ 802.1p VPT
- ▶ Type of traffic:
 - ▶ Unicast
 - ▶ Multicast
 - ▶ Broadcast
 - ▶ Known
 - ▶ Unknown



Traffic classification & actions performed in hardware without performance-degradation

Multifield classifier helps in identifying all significant traffic anchors that can be used for policy rules. The actions from the abovementioned classification can be:

Accept the traffic flow

Forward traffic that matches classification criteria.

Forwarding action can be based on destination MAC address or destination virtual interface.

Drop the traffic flow


Filter traffic matching the classifier - helps to enforce security policies by removing hostile traffic that could potentially compromise the overall network behavior.

Mark Service Level

Traffic in each traffic class for a subscriber/application can be placed in a different hardware queue for QoS scheduling.

Mark/remark VPT/DSCP/MPLS layers

Change the ToS/DSCP/DiffServ settings in the IP header, Ethernet 802.1p or MPLS EXP settings. The values can be set to trusted/mapped mode by the network elements in the interior (aggregation) network with per flow QoS processing.



Marking (setting) QoS policy in packet header

Mirror to CPU for troubleshooting or sniffing

OS9000 incorporates an internal sniffer that is used to troubleshoot network applications with current traffic to the CPU via a specific interface. The traffic can be stored into a log and inspected with internal 'grep' Linux-based tool.

Redirect to CPU for specific manipulation and alternate rule

Police to Single rate 3 colors / Police to Dual rate 3 colors (CIR/PIR)

Ingoing traffic is policed with granular bandwidth rates to enforce traffic contracts, and to specify how much traffic a specific subscriber/application can send over the provider's network.

The police action can be applied to an entire interface, a subscriber's VLAN or classes that match classified criteria.

Conformed and exceeded (out-of-profile) traffic can be assigned a color to signify discard priority and congestion avoidance with WRED on color significance.

Coloring helps to mark the traffic green, yellow or red to indicate discard priority in hardware queues. The coloring method allows bandwidth profiles' flexibility as follows:

1. Conforming traffic can be classified as green and exceeding traffic will be red, thus dropped.
2. Conforming traffic can be classified as green and exceeding traffic will be yellow, thus transmitted if no congestion occurs.
3. Conforming traffic can be classified as green and exceeding traffic will be reclassified as yellow and sent to the next statement condition.

Traffic shaping on outgoing physical interface and hardware queue

Controllable aggregated rate for bi-directional bandwidth control and service-aware QoS. Egress shaping provides fair access to multiple customers sharing and aggregating links, or aggregated traffic from multiple sites towards the subscriber/downlink.

Accounting

Traffic statistics and accounting per predefined flow that are created upon classification criteria. Accounting can be set per color, traffic type, VLAN, service level and class.

QoS at Ethernet, IP and MPLS Layers

MRV's QoS is implemented with standards-based prioritization schemes, to offer real-time packets or any other valuable service requiring a larger or prioritized share of network resources. These schemes include IEEE 802.1Q and P, the Internet Engineering Task Force (IETF) Differentiated Services (DiffServ) and the MPLS EXP bits.

Traffic at the class level is classified by class of service (CoS), IP Precedence, IP differentiated services code point (DSCP), mapped to 802.1p 3 priority bits and to provider MPLS EXP bits.

Traffic Management Services

A bandwidth profile is part of the service level specifications between provider and the customer. For premium services, traffic management requires granular hierarchies to define QoS parameters, such as Committed Information Rate (CIR) and Excesses/Peak Information Rate (EIR) from the port level down to specific customer services provisioned on the port.

The OS9000 offers MEF-compliant bandwidth profiles that can be applied at three service attributes:

- ▶ Inbound bandwidth profile at the User Network Interface (UNI)
- ▶ Inbound bandwidth profile per Ethernet Virtual Circuit (EVC)
- ▶ Inbound bandwidth profile per CoS

The bandwidth profile service consists of two offerings and can be defined in 1 Kbps increments:

1. Committed Information Rate (CIR)
2. Excess Information Rate (EIR)

CIR offers a guaranteed throughput in bits per second that the service provider commits to supplying to the customer.

EIR offers the rate (bps) above the CIR at which the traffic is allowed into UNI and that can get delivered if the network is not congested.

Both rates, CIR & EIR have burst size parameters CBS & EBS, that can be greater or equal to the largest Maximum Transmission Unit (MTU) size and expressed as bytes.

The bandwidth profile services offer effective network resources utilization and related QoS performances binding: latency, jitter and packet loss.

The OS9000 hardware architecture enables low-latency queuing and forwarding, assigned for real-time applications that will not experience latency issues that can affect performance and quality. Typical OS9000 latency is only a few microseconds, thus, the accumulated number of nodes from access through aggregation and backbone will not be over 150 milliseconds (ITU-T G.114 recommendations) along 10 msec jitter and 1% of average packet loss in an optimal-designed network.

An internal round trip reporter tool offers ongoing monitoring of performance to inspect service level end-to-end. Latency, jitter, and packet loss values, in addition to bandwidth, serve to determine whether the network flow in its present configuration can provide the requisite level of service essential for time-sensitive applications such as VoIP and video streaming.

Enhanced hard-QoS mechanism - MPLS control plane

The common DiffServ solution is best characterized as "soft-QoS" with a concept of Per Hop Behavior (PHB) that relies on an over-engineering method. As a result, the network should have a baseline of worst scenario traffic, to have predictable behavior to compensate on non-deterministic network patterns.

The MPLS technology offers more deterministic QoS, guaranteeing bandwidth in much the same way as the connection-oriented approach of ATM.

By combining MPLS technology with access list engine on OS9000 series, MRV introduces hard-QoS with end-to-end traffic engineering, ATM-like QoS signaling, path and node protection.



The OS9000's control plane leverages rich QoS signaling protocols (like RSVP-TE, IS-IS-TE) that offer traffic engineering, bandwidth reservation and admission control with an analogy to connection-oriented paths for voice, video, and other mission-critical traffic services over a provider network. The OS9000 can function as a Label Edge and Label Switch router with the capability to create traffic engineered LSPs called tunnels. These tunnels can be created using either CR-LDP (LDP tunnels) or RSVP-TE (RSVP tunnels). One of the important constraints that the administrator can define for a tunnel, is the amount of bandwidth needed for the tunnel. While the tunnel is established, the bandwidth is reserved on all the MPLS devices along the path. If according to the internal admission control there weren't enough bandwidth available on one of the MPLS path devices, that tunnel would either fail or replace an existing tunnel with lower priority. The reservation protocols are more than essential and they are enabled by the MPLS technology, which can guarantee end-to-end QoS and preserve that capability as it scales. This is a crucial function that forces the Ethernet service to be adaptable and to offer applications with carrier-class hard-QoS concept.

Conclusion

The OS9000 solution is purpose-built for triple play (video, voice, data) converged IP/MPLS/Ethernet networks. Its QoS is reinforced with the latest MEF (Metro Ethernet Forum) and IETF (Internet Engineering Task Force) standards to help service providers meet the growing demands of business and residential customers. The ability for a service provider to differentiate traffic types for a customer provides new revenue opportunities and a key differentiation in their competitive environment. The OS9000 offers the most flexible and cost-effective service aware platform for providers' first & second mile aggregation fitting today's and tomorrow's needs.

Contact:
Zeev Draer
Zdraer@mrv.com