# Differentiated Services and Integrated Services use of MPLS

Nicolas Rouhana
*University Saint-Joseph – Lebanon*
*Nicolas.Rouhana@usj.edu.lb*

Eric Horlait
*University Pierre et Marie Curie – France*
*Eric.Horlait@lip6.fr*

## Abstract

*All the new emerging QoS service architectures are motivated by the desire to improve the overall performance of IP networks. Integrated Services (Intserv), Differentiated Services (Diffserv), MultiProtocol Label Switching (MPLS) and constraint-based routing are all technologies starting to coexist together in today's Internet to provide means for the delivery of end-to-end QoS to applications over heterogeneous networks.*

*In this paper, we propose DRUM (Diffserv and RSVP/intserv Use of MPLS), an architecture that delivers end-to-end service guarantees for both Diffserv and Intserv networks, where part of the underlying technology used for IP transport is MPLS using Diffserv-like mechanisms for QoS provision. We also show how traffic engineering can ameliorate service differentiation, and illustrate how interoperability can be achieved between DRUM and neighboring Diffserv and Intserv networks.*

## 1. Introduction

In the past several years, works on QoS enabled networks led to several propositions. The Integrated Services (IntServ) architecture [1] was first introduced along with the RSVP signaling protocol [2] that applications used for setting up paths and reserving resources towards receivers before sending data. The Differentiated Services (DiffServ or DS) architecture [6], a more scalable solution, classifies packets into a small number of aggregated flows or service classes that specified a specific forwarding treatment or Per Hop Behavior (PHB). The MultiProtocol Label Switching (MPLS) [9] architecture, originally presented as a way of improving the forwarding speed of routers, is now emerging as a crucial standard technology that offers new QoS capabilities for large-scale IP networks. Furthermore, traffic engineering associated with constraint-based routing have the ability to compute routes subject to multiple constraints such as bandwidth or delay requirement, and constitute important tools used by MPLS for arranging how traffic flows through the network and improve network utilization.

All these service architectures are now viewed as complementary in the pursuit of end-to-end QoS provisioning. For example, uneven traffic distribution can be a problem for Premium service in a DS domain [8] because aggregation of Premium traffic in the core network may invalidate the assumption that the arrival rate of premium traffic is below the service rate and Differentiated Services alone cannot solve this problem; traffic engineering is needed to avoid congestion.

Figure 1 shows how these different technologies would fit together in today's Internet: an MPLS core network consisting of Label Switch Routers (LSRs) providing transport and QoS guarantees for boundary "customer" networks supporting Intserv and Diffserv architectures.

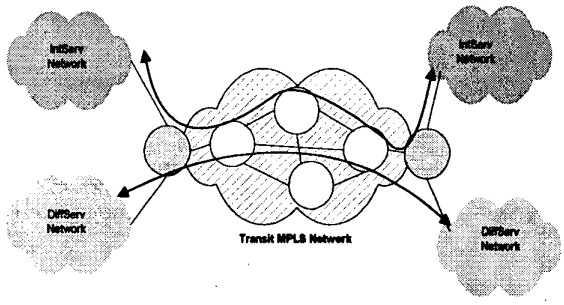


Figure 1. End-to-end QoS network

The QoS model depicted in figure 1 considers that senders from the Intserv (resp. Diffserv) networks need to communicate with receivers in other Intserv (resp. Diffserv) networks through the MPLS transit network. This is consequent to the capability of MPLS to provide "Virtual Private Networks" (VPNs) to connect organizations to their multiple sites with compatible end-to-end QoS needs.

The DRUM architecture described in this paper proposes simple mechanisms for MPLS edge and core LSRs to deliver end-to-end service guarantees and scalable QoS for both Diffserv and Intserv neighboring networks. The next section introduces a service definition in DRUM. Section 3 explains the internals of the LSRs. Section 4 presents the necessary interworking functions between the different networks, and section 5 analyses



0-7695-0722-0/00 $10.00 © 2000 IEEE

DOCKET ALARM
Find authenticated court documents without watermarks at docketalarm.com.

simulation results.

## 2. Service support in *DRUM*

MPLS being a core technology, the focus of QoS support in MPLS networks is scalability, which is achieved by flow aggregation that ensures individual end-to-end QoS guarantees without maintaining awareness of individual flows on each segment of their path. Diffserv mechanisms are therefore a good candidate to provide QoS within MPLS networks because services are based on a per-hop model and aggregate forwarding resources (buffer space, bandwidth, scheduling policy) that are pre-allocated in the LSRs for each service type. Functions such as classification, marking and policing are only needed at the edge LSRs of the network while core LSRs need only to have PHB classification, hence the scalability at the core.

To support service differentiation in DRUM, labeled packets are divided into separate traffic classes. The inherent characteristics of MPLS and Label Switched Paths (LSPs) make it easy to support aggregated flows. When an aggregation of flows is placed inside an LSP, the result is a traffic trunk [13]. Many different trunks, each with its own traffic class, may share an LSP and the 3-bit Exp field of the MPLS packet header could then be used to indicate the service class of each packet. In that case, no more than eight Behavior Aggregates (BA) can be defined within the MPLS network. If more than 8 BAs are required, the service class should then be inferred from both the MPLS label and the Exp fields. This latter scheme yields in less scalability than the former [10] and is not considered in our case.

DRUM proposes the following service classes, inspired from [7]:

1. A Gold class, consisting of a low loss, low latency and low jitter service for delay-sensitive traffic. The network commits to deliver user datagrams at a rate of a Peak Data Rate (PDR) with minimum delay requirements. Datagrams in excess of PDR are discarded.
2. A Silver class and a Bronze class for throughput-sensitive traffic. Packets in the Silver class experience lighter load (and thus have greater probability for timely forwarding) than packets assigned to the Bronze class. Packets within each class are further separated by two drop precedence levels (high and low). Within each class, the network commits to deliver with high probability user datagrams at a rate of at least a Committed Data Rate (CDR). The user may transmit at a rate higher than CDR but datagrams in excess of CDR have a lower probability of being delivered.
3. A default Best Effort (BE) service class with no expected guarantees from the network. A Less than Best-Effort (LBE) class can also be considered for background traffic or "demoted" traffic that is out-of-profile. This latter case remains to be further studied as it re-orders flows which may be undesirable.

The mapping between the Exp-field values and the BA are defined by the network operator and are MPLS network specific. Table 1 proposes mappings of the Exp field value to a pair <FCI,DPI>, where the FCI (Forwarding Class Indicator) value indicates an MPLS forwarding class and the DPI (Drop Precedence Indicator) value indicates a level of drop precedence, used by the congestion avoidance mechanisms in DRUM described later.

Table1. Exp to MPLS service class mapping

| MPLS service class | MPLS *EXP* field value FCI | DPI | Drop precedence |
|---|---|---|---|
| Network ctrl | 11 | 1 | N/A |
| Gold | 11 | 0 | N/A |
| Silver | 10 | 1 | Low |
| | 10 | 0 | High |
| Bronze | 01 | 1 | Low |
| | 01 | 0 | High |
| Best Effort | 00 | 1 | Low |
| | 00 | 0 | High (less than BE) |

The Exp values of 7 and 6 are reserved for the highest priority traffic (e.g. network control, signaling, routing updates, etc.) and Gold traffic respectively. Drop precedence levels in those classes are not defined.

## 3. The LSRs internals

Figures 2 and 3 show the internal architecture of a core LSR and an edge LSR respectively used in DRUM.

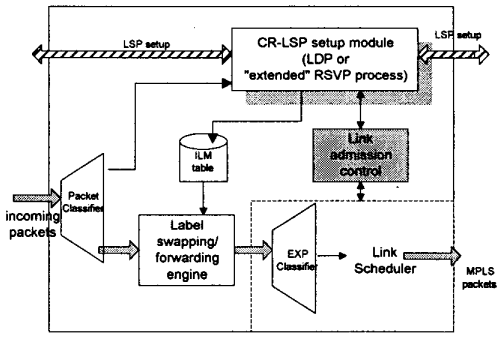


Figure 2. Functional elements of a core LSR

Both types of LSRs include functions for constraint-based routing LSP (CR-LSP) setup with link-admission control and scheduling behaviors. In addition, the ingress LSR is also responsible for classification, policing and shaping rules, LSP admission control, and interworking



DOCKET ALARM
Find authenticated court documents without watermarks at docketalarm.com.

functions (IWF) with the neighboring Diffserv and Intserv domains.

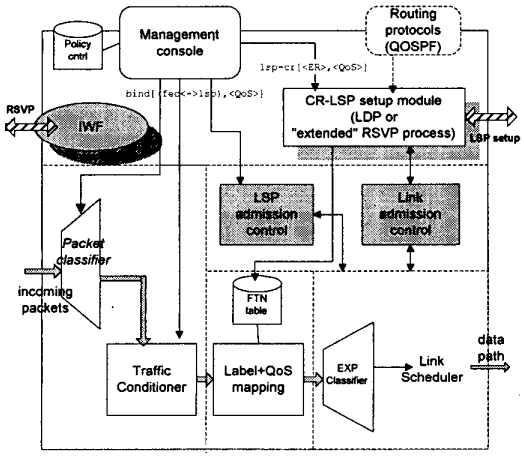


Figure3. Functional elements of an edge LSR

The interactions between these modules are detailed in the next sub-sections.

### 3.1 LSP setup and admission control

The management console at each edge LSR is used by the network manager to setup the Constraint-Based Routing LSPs by specifying the Explicit Routes (ER) and the associated Traffic Parameters, which, for the sake of simplicity, are considered in the form of token buckets values (r,b). These parameters should be chosen sufficient to accommodate the traffic of all classes to be forwarded on that LSP, i.e. must reflect at least the "sum" of all the traffic parameters of the flows to be reserved traversing the LSP. The CR-LSP setup module uses either the generic Label Distribution Protocol (LDP) [11], either an extension of RSVP [5]. Labels are allocated in a control-driven, downstream-on-demand approach which is a scheme providing more network control (e.g. all LSRs belonging to the same LSP perform the label binding in an ordered manner) and better scalability in resource conservation.

If LDP is used, the CR-LSP setup module first checks the link admission control module of its outgoing interface to the next hop in the ER to try reserving the required bandwidth. If successful, the remaining capacity of the link is diminished by (r,b) and a Label Request (LR) message is sent to the next hop in the ER of that LSP. The next hop LSR (a core LSR) also checks its link admission control to setup a reservation on its outgoing interface and so forth until the egress LSR of the ER is reached. The egress LSR then sends a Label Mapping (LM) message back to the originating LSR (OSR) – following the reverse explicit route path – with the label information that is stored in the Incoming Label Map (ILM) table within each core LSR. If the LSP setup fails due to insufficient resources along the explicit path, an error message is sent back to the OSR "tearing" down all reservations, and the administrator would then try another path. Once the LSP is setup, the desired requested bandwidth would then be available end-to-end on the explicit route for the "sum" of all aggregate traffic in all the classes.

For example, figure 4 shows an LSP between LSR1 and LSR2 of a reserved 100kbps aggregate end-to-end capacity, and an LSP between LSR1 and LSR3 of a reserved 50kbps aggregate end-to-end capacity. The remaining 1.85Mbps would still be available from the 2Mbps link on the outgoing interface of LSR1.

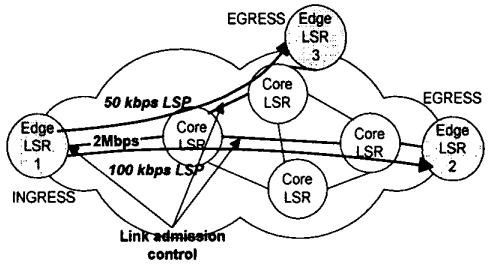


Figure 4. End-to-end link admission control

At the ingress LSR, table 2 shows the bandwidth characteristics associated with each LSP.

Table 2. Example of LSPs and their corresponding bandwidth capacity

| LSP_ID | Label out | Capacity |
|---|---|---|
| 3 | 35 | 50kbps |
| 4 | 23 | 100kbps |

The network administrator can now start allocating bandwidth statically for each service class within these LSPs, much like VPs in ATM.

### 3.2 Packet classification

Packet classification is a function required at the edge of the MPLS network. Its goal is to provide identification of the packets belonging to a traffic stream to a Forward Equivalence Class (FEC) [9]. The classifier can be a Multi-Field (MF) classifier, which performs packet selection based on the combination of one or more header fields in the incoming IP packet (source address IP and/or port, destination address IP and/or port).

Once the CR-LSPs have been setup, the next task is to configure the classifier to bind a particular flow and its traffic parameters (r,b) to an LSP and assign the flow to a particular service class. For example, if an organization wishes to reserve a certain amount of bandwidth to interconnect its two sites across an MPLS network, the



DOCKET ALARM
Find authenticated court documents without watermarks at docketalarm.com.

network manager "tunnels" the customer's flow on an already established LSP across his network with remaining characteristics satisfying the customer's QoS requirements. An LSP admission control module at the ingress LSR provides control-load support on that LSP at the ingress. This module measures whether bandwidth is still available on that LSP for the traffic parameters requested by the new flow being added to that path and stores the information in the FTN (FEC to Next Hop Label Forwarding Entry) of the edge LSR. Various admission control algorithms [12] are used in DRUM to provide control-load support on that LSP.

When a packet is received from a neighboring network, the edge LSR "encapsulates" the packet in an MPLS header with the appropriate label and Exp value. Table 3 shows the binding of two Gold flows to the same LSP that are each using 20kbps of the bandwidth of the LSP with LSP_ID 4.

Table 3. Example of an FTN table for LSP_ID 4 and its remaining capacity of 60 kbps

| FEC | LSP_ID | EXP value | Label out | Flow capacity |
|---|---|---|---|---|
| flow1 | 4 | 110 | 23 | 20kbps |
| flow2 | 4 | 110 | 23 | 20kbps |

### 3.3 Traffic conditioners

Traffic conditioners are a vital part of a differentiated services architecture as they perform the necessary policing actions on incoming packets at the edge of the network. They act on the classified packets, and, as shown in figure 5, consist of leaky buckets associated with each incoming "Gold" traffic, and a token bucket for each "Silver" and "Bronze" traffic. Packets that are out-of-profile are either discarded (in the Gold class), either given a high drop precedence (in the Silver/Bronze/BE class).

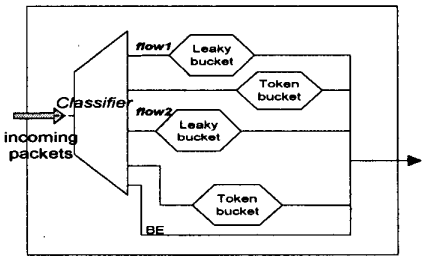


Figure 5. Traffic conditioners

### 3.4 Per-Hop scheduling classes

Several types of scheduling behaviors and drop *policies* may be used to deliver the forwarding behavior described in section 2. One simple example is given in figure 6, which consists of four different queues, one queue per traffic class, with a simple priority scheduler serving the queues. MPLS packets are classified according to the Exp field and forwarded to the appropriate queue.

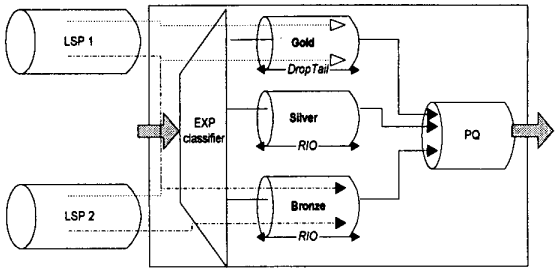


Figure 6. Output scheduler in the LSRs

Since all the LSRs support the same PHB ←→ EXP mapping, LSPs are merged implicitly, and traffic of the same class from different LSPs are statically multiplexed together in the same queue. The different admission control mechanisms and traffic conditioners described earlier protect low priority queues from being starved by the high priority queues. The Gold class is the one with the highest priority, with tail-drop discard giving the minimum service delay for the packets. Each of the other classes uses a separate queue managed by the congestion management scheme *Random Early Detection (RED) with In and Out* (RIO) [14].

## 4. Interworking with Diffserv and Intserv domains

In this section, we show how the MPLS service architecture can support QoS for Intserv and Diffserv networks and provide interoperability at the boundary. In order for a customer to receive services from the network, he must have a service level agreement (SLA) with the provider. SLAs can be static or dynamic; static SLAs are negotiated on a regular (e.g., monthly or yearly) basis, while dynamic SLAs require the use of some signaling protocol (e.g., RSVP) to request service on demand. These dynamic requests are taken into account by the IWF module at the ingress LSR, which "tunnels" them through the MPLS network to the destination, thus ensuring that flow reservation happens end-to-end [13]. Also, in order to support both types of SLAs and minimize the LSP setup delay, the Service Provider can provision his network by statically allocating the necessary resources and setting up all the constraint-based LSPs between the MPLS end-points, based on customers' needs and anticipated incoming traffic patterns through the SLAs.



DOCKET ALARM
Find authenticated court documents without watermarks at docketalarm.com.

Ingress LSRs map the incoming QoS requests from the neighboring Diffserv and Intserv networks to the corresponding service class within the MPLS network. In this model, each incoming flow is assigned to one of the available classes for the duration of the flow and traverses the MPLS cloud in this class. The service mappings given in table 4 follow most naturally from the service definitions: Guaranteed Service [3] and Expedited Forwarding [8] map to the Gold MPLS class, and Controlled-Load service [4] and Assured Forwarding [7] map to Silver and Bronze classes.

Table 4. Service map from Intserv/Diffserv to MPLS

| Incoming service requirements IntServ | DiffServ | MPLS matching service |
|---|---|---|
| Guaranteed Service (GS) | Expedited Forwarding (EF) | Gold |
| Controlled-Load (CL) | Assured Forwarding (AF) | Silver/Bronze |
| Best Effort | Best Effort | Best Effort |

For completeness, we propose in table 5 possible mappings for the service combinations and identify how the Exp field can be used in the MPLS header of packets across the MPLS network to obtain the equivalent service. Taking into account that the 8 BAs defined within the MPLS network offer less service granularity than the Diffserv classes (one EF and four AF with three possible drop precedence levels in each class [7]), the network administrator can choose to group together traffic flows requiring similar service into a single MPLS service class, and configures the mappings of the DSCP values that Diffserv QoS uses, as well as the incoming Intserv service request, to the Exp value for output port scheduling and congestion avoidance mechanisms.

Table 5. Proposed QoS mappings

| Intserv service type | Diffserv class PHB | DSCP | MPLS EXP field | MPLS service class |
|---|---|---|---|---|
| GS | EF | 101110 | 110 | Gold |
| Control load | AF11 | 001010 | 101 | Silver |
| | AF12 | 001100 | 100 | |
| | AF13 | 001110 | 100 | |
| | AF21 | 010010 | 101 | |
| | AF22 | 010100 | 100 | |
| | AF23 | 010110 | 100 | |
| | AF31 | 011010 | 011 | Bronze |
| | AF32 | 011100 | 010 | |
| | AF33 | 011110 | 010 | |
| | AF41 | 100010 | 011 | |
| | AF42 | 100100 | 010 | |
| | AF43 | 100110 | 010 | |
| BE | DF | 000000 | 000 | Best Effort |

For instance, in table 5, one choice might be to give the classes AFx2 and AFx3 (x=1,2,3,4) the same drop precedence within each service class 'x'. For Controlled-Load service requests incoming from Intserv networks, the Exp value can vary between Silver or Bronze service depending on the network provider's pre-configured settings based on per customer criteria.

## 5. Validation

In this section, we briefly describe a simulation setup used to validate DRUM in providing service and delay differentiation for Intserv and Diffserv networks. We used the network simulator ns-2 [18] for our analysis with new modules supporting the DRUM architecture described earlier. The sample network topology is given in figure 7, and consists of seven edge LSRs (n1, n6, n7, n8, n9, n10), four core LSRs (n2, n3, n4, n5) and links at 2 Mbps. Three classes of traffic were used: Gold traffic at CBR rate of 300 kbps and 600 bytes packet sizes, Silver traffic at CBR rate of 200 kbps and 500 bytes packet size, and Best-Effort traffic at 500 kbps CBR rate with 1000 bytes packet size. Nodes n1 and n6 each generates one Gold, one Silver and one Best-Effort flow towards nodes n7 and n10 respectively, whilst node 8 sends to node 9 one Gold and one Best-Effort flow.

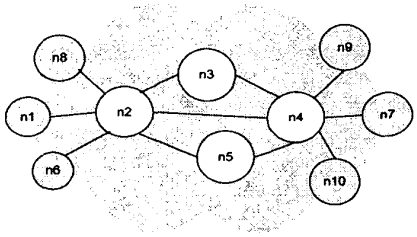


Figure 7. Sample network topology

A first test case was to mix all the flows on the same path n2-n3-n4. Figure 8 shows rate guarantees for the Gold and Silver traffic at node 10 with no loss since the flows send at the subscribed rate, while the Best Effort traffic used up the remaining bandwidth and was subject to packet loss.

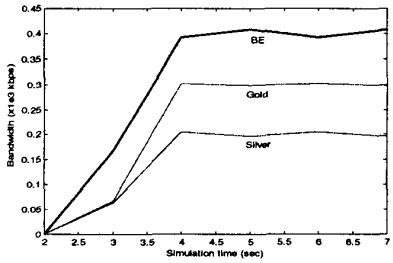


Figure 8. Throughput characteristics



DOCKET ALARM
Find authenticated court documents without watermarks at docketalarm.com.

DOCKET ALARM

# Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

## Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

## Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

## Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

## API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

**LAW FIRMS**
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

**FINANCIAL INSTITUTIONS**
Litigation and bankruptcy checks for companies and debtors.

**E-DISCOVERY AND LEGAL VENDORS**
Sync your system to PACER to automate legal marketing.