Routing architecture in DiffServ MPLS networks

Gonzalo Camarillo
Advanced Signalling Research Laboratory
Ericsson, FIN-02420 Jorvas, Finland
Gonzalo.Camarillo@ericsson.com

Abstract

The Internet is currently requiring a means for providing different users with different service levels. Traffic generated by users with high requirements has to be treated with certain priority while the traditional best effort services are still made available.

Differentiated Services enabled networks classify the traffic into different priority levels then apply different scheduling and queuing mechanisms at the nodes. MPLS allows traffic engineering and efficient routing of packets combining layer 2 and layer 3 mechanisms.

In order to combine these two mechanisms a way is needed to distribute labels between the MPLS routers. The MPLS routers provide information about the service level associated with the traffic. MPLS labels also contain typical routing information. LDP with some new extensions allows this functionality.

1 Introduction

The Internet is evolving from a network that provided just best-effort transportation to a network capable of providing a wide range of services. The delivery of data related to some services has tougher requirements than best-effort packets. This has to be differentiated at network level

There have been some different proposals to enhance the Internet in order to support the different requirements of different types of traffic. The Integrated services architecture [1], with RSVP [2] as main protocol has severe scalability problems. It has been proven to be unsuitable for large networks since every single micro flow is treated individually. Therefore, this architecture cannot be implemented in the core of the Internet where a large number of flows coexist.

The DiffServ (Differentiated Services) architecture [3] pushes the complexity to the edge of the network where there are fewer flows in parallel. Traffic classification and packet filtering is performed here. Micro flows are aggregated into traffic classes, solving the scalability problem in the core of the network.

The MPLS architecture [4] represents a similar approach. MPLS also follows the IP philosophy and places the complexity at the edge of the network. Interior routers are simpler and need less functionality than the ones located at the edge.

It is possible to combine features of DiffServ and MPLS in the same network. Thus, the users can take advantage of getting a better service.

2 Differentiated Services

The basic goal of Differentiated Services architecture [3] is to fulfill the performance requirements of the users. Users request a certain performance level and the network provides it as long as the user traffic has certain characteristics. The performance level provided and the characteristics of the traffic to be injected in the network are defined in a SLA (Service Level Agreement) [5].

The part of the SLA dealing with technical details is referred to as SLS (Service Level Specification). Inside the SLS, the TCS (Traffic Conditioning Specification) specifies the expected performance (throughput, drop probability, latency…), the profile of the traffic to be used (peak data rate, burst size…) and actions to perform in case of excess traffic.

Excess traffic is typically assigned a higher drop probability, it is delayed or simply discarded. There are routers at the edge of the network that measure the traffic and reshape it so that it falls inside the TCS. These edge routers also classify the data packets into several different traffic classes.

These traffic classes have a service level associated with them. Once the packets are classified and tagged at the edge of the network they are treated accordingly inside the network.

Different traffic classes have different priority levels in the routers' queues. Scheduling algorithms have to ensure high priority packets are forwarded before low priority ones. These algorithms must also ensure a certain minimum bandwidth for best-effort traffic.

2.1 Per-Hop Behavior

The way each router in the path treats a packet is referred to as PHB (Per-Hop Behavior). The default PHB is best-effort but there are more PHBs already defined or being standardized. There are also PHBs that are locally defined inside a node and do not correspond to a well known set of features. All PHBs are local to the node implementing them.

The definition of a PHB does not include the specific algorithm to be employed. It outlines a set of requirements that have to be fulfilled in order to provide that specific PHB. Examples of PHB are EF (Expedited Forwarding) and AF (Assured Forwarding).

EF [6] requires the departure rate (from the node) to be equal or greater than a configurable rate. The time for measuring this rate is the time interval equal to or longer than the time it takes to send an output link MTU packet at the configured rate.

Traffic to which EF PHB applies is forwarded as soon as possible independently of the state of the node. EF traffic is not delayed in queues as far as it is possible. EF PHB is used for very high priority traffic.

EF PHB can be implemented using several algorithms such as CBQ or single priority queues. Different algorithms still fulfilling the EF requirements introduce different jitter to the packets.

AF [7] is a PHB group that provides packet delivery in four independent forwarded classes with three levels of drop precedence. Packets that belong to a class are not reordered in the node and nodes do not aggregate two classes together.

A minimum amount of forwarding resources has to be allocated for each class, and available resources in the node can be statistically allocated to these classes also. The three levels of drop precedence have to be implemented using at least two different loss probabilities.

The naming convention used is AFij, where i represents the class and j the drop precedence l
evel. The higher j is the more probable the packet will get discarded.

AF PHB can be implemented in different ways. RED [8] (Random Early Detection) is recommended to signal congestion to the end points.

2.2 Differentiated Services Code Points

Once PHBs have been defined for individual nodes there is a need of providing coherent treatment to the packets along the whole path. Every node traversed by packets that belong to a certain class has to apply the same PHB to the packets. This is achieved by adding a tag to the packets which describes the PHB required.

This information is placed in the so-called DS (Differentiated Services) field. The DS field supersedes the existing definitions of TOS (Type of Service) in IPv4 [9] and the traffic class octet in IPv6 [10].

The TOS field contains eight bits. Six of them are used to represent the DSCP (Differentiated Services Code Point). Two bits remain unused.

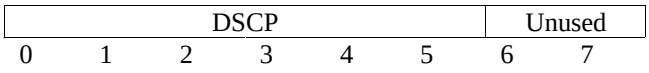

Figure 1 : DSCP format

The default best-effort service corresponds to the DSCP equal to zero (DSCP='000000'). There are some reserved code points of the form 'xxx000' called class selector code points. Some routers currently use these code values. Therefore, PHBs associated to class selector code points must fulfill some requisites (for example, the higher the code point is the higher the priority the packet has). This provides backward compatibility with old implementations.

Apart from class selector code points the DSCP space is divided into three pools of code points. The first pool of code points ('xxxxx0') contains just standardized values. Pool number two ('xxxx11') is intended for experimental and local use. The third pool ('xxxx01') is a mixture of standardized and experimental point codes. It will be used for standard point codes when/if pool number one runs out of available DSCPs.

The standard DSCP for the EF PHB is '101110'. The table below shows the standard code points for the AF PHB group.

Table 1: AF code points

	Class 1	**Class 2**	**Class 3**	**Class 4**
Low drop precedence	001010	010010	011010	100010
Medium drop precedence	001100	010100	011100	100100
High drop precedence	001110	010110	011110	100110

Mapping between DSCPs is needed when non-universal DSCPs are used. The edge routers between domains perform this mapping.

2.3 Services provided

Making use of the DSCPs and the respective PHBs triggered in the nodes it is possible to provide the users with differentiated services. Inside the SLA, the user accepts to send data according to the TCS, and the network compromises to provide a certain level of service.

The edge of the network has to undertake the classification of the flows and check if they are inside the TCS. A set of devices are implemented in the edge routers for this purpose. They are classifiers, meters, markers, droppers and shapers. BA (Behavior Aggregated) classifiers classify the packets based solely on the DS field while MF (Multiple Fields) classifiers are based on multiple fields. MF classifiers can provide per-flow services.

Meters measure the incoming traffic paying attention to the parameters that appear in the TCS. Depending on the applicable actions, excess traffic or not conforming to the TCS is passed to a marker, a shaper or a dropper.

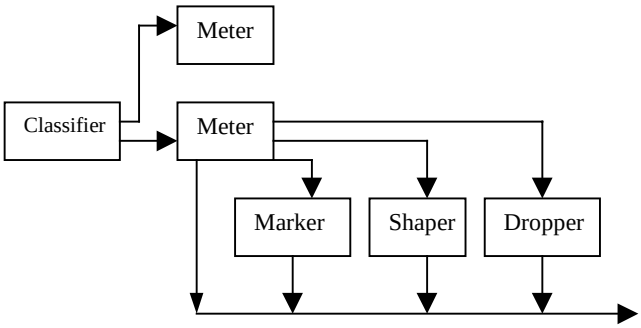

Figure 2: Structure of an edge device

A marker remarks the traffic with a different DSCP. This implies this traffic will get a different PHB in the network. Shapers delay the traffic so that traffic in the network conforms to the TCS. Droppers simply discard the packets.

A service can be defined in different ways. Quantitative services are defined using concrete figures (i.e. 80 % of the traffic will experiment less than 30ms latency). Qualitative services are defined in a more loose manner (i.e. traffic will experiment low latency and low jitter). There is a third possibility to define a service. It is referred to as relative quantification service. It consists of comparing the performance of two different classes (i.e. traffic at level A will have double allocated bandwidth than traffic at level C).

Quantitative service traffic has always higher priority than qualitative traffic. The network should always have resources available for this kind of traffic. However, a certain ISP (Internet Service Provider) can design a network taking into consideration the statistical gain. That is, since most likely all the quantitative sources will not transmit at the same time the network can be oversubscribed.

A service is defined within its scope. The scope is usually based on the ingress point. It can apply to all the egress points or just a set of them. SLAs based on destination points are difficult to implement. Destination addresses are not typically permanent (when DHCP is used) and denial of service attacks are possible (i.e. the attacker overloads the destination point with traffic). SLAs based on destination address can also have collisions with the SLA of the sender. SLAs based on destination points are usually applied on a micro-flow basis.

Table 2 : TCS example

DS mark	Profile	Scope	Non-conforming
AF11	1 Mbps	Any egress point	Marked as AF13

The table above shows an example of TCS. All the traffic up to 1 Mbps addressed to any destination will be marked as AF11. The excess traffic will be marked as AF13, it will belong to the same class but with higher drop probability.

3 Multi Protocol Label Switching

MPLS [4] (Multi Protocol Label Switching) consists of combining layer 3 routing with layer 2 switching. It is not confined to any specific link layer. However MPLS can take advantage of some services provided by some specific link layers such as ATM or Frame Relay.

MPLS provides connection oriented switching. Labels are associated to streams of data. Packets belonging to those streams are forwarded based on their labels.

Labels are short fixed length, physically contiguous locally significant identifiers used to identify a stream. Routers exchange labels and the information associated to the labels between them. They can use a dedicated exchange protocol or piggyback this information in another protocol already in used in the network.

MPLS presents some advantages [11] over a network that employs just layer 3 routers.

MPLS employs a simpler forwarding paradigm. MPLS forwarding is based on exact matches of labels. Layer 3

routing is based on comparisons of CIDR addresses. The matching rules are more complex and the IP addresses longer than MPLS labels.

MPLS provides a feasible way of implementing traffic engineering with different granularity levels. Explicit routing can be performed without carrying the whole explicit route in every single IP packet transported. The packets are labeled by the ingress edge MPLS router and are routed accordingly.

MPLS moves heavy processing to the edge of the network. Packets are classified and labeled by the edge routers. Interior routers have just to perform a label look up, swap the label to one meaningful to the next hop and decrement, if applicable, the TTL (Time To Live) of the packet.

MPLS also presents some advantages when compared to a network based on ATM or Frame Relay. The scalability is improved due to the reduction of logical links between the switches. An LSR (Label Switching Router) does not have to maintain logical connections with all the LSRs in the network.

A smooth operation over different link layer technologies is achieved. ATM based networks can operate together with networks based on Frame Relay or with PPP links.

3.1 Routing in MPLS networks

A packet enters into an MPLS network through an edge router. Edge routers have a number of FECs (Forward Equivalence Class). A FEC defines which packets have to be forwarded in the same way. A FEC can consist of a destination address or it can have a finer granularity (i.e. source address, destination address and both port numbers).

Once the packet is classified into one FEC it is assigned a label. This label is meaningful just locally, between the edge router and its LSR peer which is the next hop for the packet (according to the layer 3 routing protocols or to traffic engineering if it is performed).

The packet is forwarded to the next LSR. The LSR looks up the label contained in the packet and swap it with the label associated to this FEC that is meaningful to the next peer LSR. This process continues until the packet reaches the edge egress router or its destination.

Table 3: Label table in an LSR

Incoming		Outgoing	
Port	**Label**	**Port**	**Label**
1	A	3	C
2	B	4	D

MPLS uses layer 3 protocols for discovering the next hop in the path for the destination.. The label just makes it simpler to find a match in the router's routing table for the forwarding decision.

Packets can contain several labels. They are placed in a label stack. This can be used when a routing hierarchy exists. The first label can be utilized to forward the packet inside an administrative domain. Thus, it will be based on the information provided by interior routing protocols. The second label can be meaningful between BGP (Border Gateway Protocol) peers. This label will be based in the exterior routing protocol used between both domains (BGP in this example).

3.2 Label assignment and distribution

MPLS is not tied to any concrete label distribution protocol. This section describes different features of these protocols. Every label distribution protocol implements some of them and every concrete network is configured using different modes of operation.

There might be several events that trigger label assignment and distribution in a network. There are basically associated with the arrival of control traffic or data traffic.

Control traffic is generated due to topology changes (routing protocols) or due to path reservations in the network (RSVP-like protocols).

The arrival of data traffic can also trigger the distribution of new labels that make it possible to route the new incoming stream.

Topology driven label assignment is always used in MPLS networks. However, the following sections show that when MPLS is used in conjunction with DiffServ the other methods have to be used as well.

Topology changes trigger control traffic that informs the routers about the new topology. When the topology changes new labels are distributed.

Labels are used between peer LSRs. There are two types of label allocations depending on which LSR chooses the actual label. In downstream label allocation the downstream LSR chooses the label to be used. Since this label is used to perform the look up in the label table in the LSR downstream it can be chosen in order to optimize this look up process.

When the LSR upstream chooses the value of the label the allocation method is known as upstream label

allocation. This allocation method can be used in multicast environments where several identical packets are sent through different ports. The upstream LSR can choose labels with the same value for all the packets, improving the performance of the forwarding process.

Currently just downstream label allocation is used in MPLS networks. MPLS in multicast environments has not yet been developed.

There is downstream independent label distribution and downstream ordered label distribution. In downstream independent label distribution an LSR sends a new label to the LSR upstream even if it does not have a label for that FEC from downstream.

In downstream ordered label distribution an LSR does not advertise a new label to the upstream LSR until a label for that FEC arrives from downstream. Consequently, the egress edge router is the first router advertising a label.

There is still a further classification. In downstream on demand label distribution an LSR requests from its next hop a label for a certain FEC. In unsolicited downstream label distribution the labels are distributed without requests.

Upon reception of a label the LSR can behave in two different ways. If the LSR is configured for using liberal retention mode it will keep all the labels received. In conservative retention mode just labels from the next hop for the FEC are accepted.

Liberal retention mode requires more resources from the LSR. It has to store labels that are not currently in use because they do not correspond to the next hop for the traffic.

Liberal retention mode has also some advantages. In case of topology changes when the next hop for a certain flow changes there is no need to request a new label. The label is already in the LSR label table and it is ready to be used by the next hop.

Label distribution protocols have to avoid packets being routed in loops in the network. There are basically three methods for avoiding loops: loop survival, loop detection and loop prevention.

Loop survival consists of limiting the time a packet can be in the network. It is usually implemented with hop counters. When the counter reaches a certain limit the packet is discarded.

Loop detection mechanisms look for loops and remove them from the network. These mechanisms do not prevent the network from creating loops. They just eliminate them when the loops are detected.

Loop prevention algorithms do not create paths that contain loops. Every path is analyzed before becoming active.

3.3 Label Distribution Protocol

LDP (Label Distribution Protocol) in one of the protocols that can be used for label distribution. LDP was developed for performing specifically this task. This section describes its functionality and its characteristics.

LDP specifies a set of procedures and messages by which LSRs establish LSPs through a network by mapping network layer information to data-link layer switched paths. LDP uses 4 types of messages: discovery messages, session messages, advertisement messages and notification messages.

Discovery messages are used to announce and maintain the presence of an LSR. They are used for discovering new LSRs and for checking if the peer LSRs to which the LSR have logical connections are up and running. This is performed via a keep alive mechanism. If the peer LSR does not respond to these messages the logical connection towards it is closed.

Discovery messages can be used between LSRs which are connected at link level or which are not. This is the case when two border BGP routers use labels between them. They are not typically connected at link level but they still have to exchange label information.

Once a peer LSR is discovered a TCP connection is established with it and session initiation is undertaken. All the parameters needed for the LDP session are negotiated through session messages. Upon completion of this negotiation the two LSR peers begin creating and exchanging label mappings for the FECs. Advertisement messages are used to establish new labels.

Notification messages are used to carry advisory information and to notify error conditions. When a procedure is not completed successfully a notification is triggered.

A FEC specification is provided for each LSP. This FEC specifies the packets that should be labeled with the label associated to the LSP in every node. LDP just provides host addresses and address prefixes for defining FECs. It will be seen that this is not enough for describing paths that require a certain service level. This FEC definition has to be enhanced in order to achieve more functionality.

LDP uses downstream label allocation. The LSR downstream chooses the actual value of the label to be used for a certain LSP. Label release is always performed by the upstream LSR. This avoids routing packets based on a label that has been or is being released by the downstream LSR.

LDP uses TLV (Type-Length-Value) encoded objects. All the information inside LDP messages is inside TLV objects. The protocol can be extended with new TLVs or enhanced definition of already existing ones. This encoding scheme makes the protocol modular and easy to upgrade.

LDP uses path vectors and hop counters for avoiding loops. The path vector consists of a vector carrying the LSRs identifications of all the LSRs in the path. Loops are detected when an LSR receives a message with its identification inside the path vector.

Hop counters are used to limit the amount of resources that can be consumed when loops are created. When the hop counter reaches a certain threshold the system acts as if a loop had been encountered in the path vector. The LSP is not established.

Labels are stored in the LIB (Label Information Base) of the LSRs. When due to transient conditions a packet is received containing a label that does not have a match in the LIB it is not safe to forward further the packet. That packet has to be discarded. Forwarding a packet based on layer 3 routing could create loops that would remain undetected.

3.4 Label format

MPLS uses different label encapsulation depending on the link layer below. In ATM, for instance, MPLS labels are the VPI (Virtual Path Identifier) and the VCI (Virtual Channel Identifier) in the ATM headers.

When the link layer does not provide a means for encapsulating MPLS labels the so-called shim header is used. It is added after the link layer header and its format is shown in the figure below.

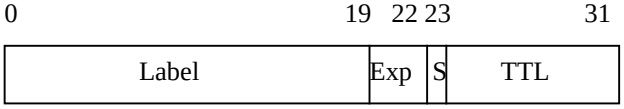

Exp: Experimental
S: Bottom of the stack
TTL: Time To Live

Figure 3: Shim header format

The shim header provides a TTL (Time To Live) field. When an LSR swaps a label encapsulated using the shim header into a label using ATM encapsulation this information is lost.

Multiple shim headers can be carried in a packet in a header stack. This is used when there is a routing hierarchy in the network.

Some values of the labels are reserved. They can indicate that the IP header after the shim header has to be examined or that the routing has to be performed based on a label different than the first one.

The shim header has 3 bits reserved for experimental and future use. Some approaches to combine MPLS with DiffServ utilize this experimental field to encapsulate information related to the PHB expected.

4 DiffServ MPLS networks

As previously stated, DiffServ and MPLS have some things in common. Both approaches push the complexity to the edge of the network. The edge routers have to deal with less flows than routers in the core. Therefore, they can perform classification of flows without being overwhelmed with the traffic traversing them.

Both MPLS and DiffServ label the packets after classifying them. Short fixed length labels are used in both networks. They are called MPLS labels in MPLS networks and DSCPs in DiffServ networks. Routers in the core network treat the packets according to these tags.

The DSCP of a packet determines the behavior of the nodes regarding scheduling mechanisms and queuing management. They typically define the priority and the drop precedence of the packets.

The MPLS label of a packet determines the path the packet takes. The packet is routed based on its label. Traffic engineering can be performed by assigning certain labels to paths with certain characteristics.

It is desirable to combine features from both MPLS and DiffServ [12] in order to achieve a good QoS in the network. Combining both approaches, it is possible to specify the paths the packets take and their behavior in the queues of different routers.

MPLS routers do not examine the IP header of the packets. However, the DSCP of the IP datagrams is contained in the IP header. Therefore, in order to use DiffServ in MPLS networks the DiffServ related information contained in the IP header has to be mapped to the MPLS label assigned to the packet.

This mapping can be done in several ways. The network administrator has to decide how the DiffServ BAs map into the MPLS LSPs (Label Switched Path).

If multiple BAs are mapped to the same LSP all the packets will have the same label. The experimental field in the MPLS header is used to specify the PHB applicable to each packet. The PHB includes scheduling and drop preference parameters. This way of mapping different BAs into a LSP using the experimental field to define the PHB is referred to as E-LSP (Exp-inferred).

The shim header provides a 3 bit experimental field. When MPLS is used with ATM, the CLP is used. In frame relay the DE bit is used. The short length of these fields limits its utilization when multiple BAs exist in a network.

If just a single BA is mapped into a single LSP, the DSCP is encoded implicitly in the label. Thus, the experimental field can be used for encoding the drop precedence of the packets. This is known as L-LSP (label only inferred). L-LSP overcomes the problem of having a short field for encoding the DSCP. However it imposes higher requirements on the system. The number of labels grows and so does the amount of resources needed in the network.

This can lead to a severe scalability problem. In a network with three different PHBs available the number of labels needed is multiplied by three. There will be three different labels indicating the same LSP but with different priority levels for the traffic. Maintaining that amount of labels can become a problem soon if the number of PHBs in the system increases.

There always is a best effort label for any LSP. This best effort label is always present and the routers change it in response to topology changes. Topology changes are noticed by the routers when the routing protocols advertise new available paths. As described in previous sections this is known as control traffic driven label distribution.

In the labels corresponding to the same LSP but with different PHBs are distributed in response to routing protocols traffic the number of labels to be maintained in the system is very large.

Another possibility inside control driven label distribution is to establish the labels when they are about to be used. The need for a LSP with a concrete PHB is signaled via a reservation protocol like RSVP. Upon reception of the request a label is established in the MPLS network.

This method reduces the number of labels to be maintained but introduces more traffic and delay in the connection. Before the traffic can be routed the LSP path has to be established.

Traffic driven label distribution does not introduce the overhead traffic produced by a reservation protocol. Labels are established upon reception of data traffic. This maintains the number of labels at a minimum, but it has also some disadvantages. The first packets of every flow cannot be provided with the QoS desired since the proper labels are being established. Thus, the beginning of every flow will not conform to the SLA between the user and the network. This can be a small drawback in networks that carry long flows. If the flow is held for a long time the time employed to establish the LSP is negligible. However, networks handling short duration flows cannot use this data traffic driven approach.

Networks can reach a trade-off and use a combination of all these techniques. Some ISPs have dial in customers as their main source of income. So, they have many intermittent connections to their network. In this situation labels corresponding to certain users that are not attached to the network at some point of time can be released. These labels can be established once the user is connected again. This reduces the amount of labels that the system has to keep track of. If there are no users connected using a certain level of service all the labels related to that service cannot be used. Therefore, the system does not need them at that point of time.

All this information can be found in the policy server. The policy server has information about the SLAs between the users and the ISP and knows which users are allowed to use certain service levels and which ones are not.

4.1 Label distribution in DiffServ MPLS networks

There are several protocols available to undertake the task of distributing labels in a Diffserv MPLS environment. Protocols that are already in use in the network can be piggybacked with extra information in order to perform this task. Another solution is to implement a protocol with the specific task of distributing labels, like LDP. This section analyzes both approaches.

There have been defined some extensions to RSVP for traffic engineering and LSP establishment [13]. RSVP establishes LSP tunnels in the network. If a network already uses RSVP, this approach can avoid implementing a new protocol. This will save network resources. It also reduces the traffic in the network since RSVP messages perform both reservation of resources and establishment of LSPs at the same time. However, RSVP also allows to establish LSP paths without QoS

requirements. Labels are associated to RSVP flows because RSVP allows very flexible definitions of flows.

DiffServ networks are typically like black clouds for RSVP. That is, RSVP is used from the end user to the border ingress DiffServ router. The next RSVP router in the egress DiffServ router. Typically there is no RSVP reservation of resources inside DiffServ networks. This implies that just the border routers need to implement it. Thus, the advantage of being an already implemented protocol in the network can be argued.

BGP [14] can also be used for exchanging and distributing labels. Label mapping information is piggybacked in BGP messages advertising new routes. This protocol can be used between border routers that are BGP capable. Therefore, it does not solve the problem of label distribution inside the DiffServ MPLS domain.

As it was seen in previous sections LDP is a protocol specifically designed for distributing MPLs labels between LSRs. With some extensions LDP can be used also for establishing LSPs with QoS requirements [15].

When LDP is employed with these extensions it is known as CR-LDP (Constraint-based Routed LDP). CR-LDP allows performing traffic engineering in the MPLS network and defining the traffic characteristics of a LSP. Some new TLVs are defined for these purposes.

CR-LDP describes LSPs with QoS requirements with a set of parameters. These parameters are the weight peak data rate, peak bust size, committed data rate, committed burst size and excess burst size. The weight parameter indicates the share of the bandwidth in excess about the committed data rate that the LSPs will get.

CR-LDP allows specifying the route that labeled packets have to take. This makes it possible for traffic belonging to different applications with different service levels to traverse a different set of nodes.

5 Conclusions

As it has been described in this document MPLS and DiffServ work well together due to some synergies between them. Both push the complexity of the network to the edge. Therefore MPLS and DiffServ edge routers perform a similar set of functions that can be combined when both are implemented in a network.

MPLS and DiffServ are complementary solutions for the problem of providing different service levels in a single network. DiffServ defines different behaviors in the nodes while MPLS deals with the paths between different nodes.

There are some scalability problems derived from the use of different labels associated to different QoS for the same path. They can be minimized without allocating all the possible labels at the same time. Different methods are available but the philosophy of all of them is not to allocate resources when no traffic is going to use them.

There are two main protocols that can be used for exchanging MPLS labels with different service levels: RSVP with extensions and CR-LDP.

Both RSVP and CR-LDP provide capability to control the nodes that the LSP passes. Both protocols are suitable for label exchange and distribution in DiffServ MPLS domains but RSVP can be more easily employed if the network is already using it for resource reservation.

6 Acronyms

AF: Assured Forwarding
ATM: Asynchronous Transfer Mode
BA: Behavior Aggregated
BGP: Border Gateway Protocol
CIDR: Classless Inter Domain Routing
CLP: Cell Loss Priority
CQB: Class Based Queuing
CR-LDP: Constraint-based Routed LDP
DE: Discard Eligibility
DHCP: Dynamic Host Configuration Protocol
DiffServ: Differentiated Services
DSCP: Differentiated Services Code Point
EF: Expedited Forwarding
E-LSP: Experimental inferred LSP
FEC: Forward Equivalence Class
IP: Internet Protocol
ISP: Internet Service Provider
LDP: Label Distribution Protocol
LIB: Label Information Base
L-LSP: Label only inferred LSP
LSP: Label Switched Path
LSR: Label Switched Router
MF: Multiple Fields
MPLS: Multi Protocol Label Switching
MTU: Maximum Transmission Unit
PHB: Per-Hop Behavior
PPP: Point-to-point Protocol
QoS: Quality of Service
RED: Random Early Detection
RSVP: ReSerVation Protocol
SLA Service Level Agreement
SLS: Service Level Specification
TCP: Transmission Control Protocol
TCT: Traffic Conditioning Specification
TLV: Time-Length-Value
TOS: Type of Service
TTL: Time To Live

VCI: Virtual Channel Identifier
VPI: Virtual Path Identifier

References

[1] Braden R., Clark D., Shenker S., “Integrated Services in the Internet Architecture: an Overview”, RFC 1633. June 1994.

[2] Braden R., Zhang L., Berson S., Herzog S., Jamin S., “Resource ReSerVation Protocol (RSVP)”, RFC 2205. September 1997.

[3] Blake S., Black D., Carlson M., Davies E., Wang Z., Weiss W., “An Architecture for Differentiated Services”, RFC 2475. December 1998

[4] Rosen E., Viswanathan A., Callon R., “Multiprotocol Label Switching Architecture”, draft-ietf-mpls-arch-06.txt, work in progress. August 1999.

[5] Bernet Y., Binder J., Blake S., Carlson M., Carpenter B., Keshav S., Davies E., Ohlman B., Verma D., Wang Z., Weiss W., “A Framework for Differentiated Services”, draft-ietf-diffserv-framework-02.txt, work in progress. February 1999

[6] Jacobson V., Nichols K., Poduri K., “An Expedited Forwarding PHB”, RFC 2598. June 1999.

[7] Heinanen J., Baker F., Weiss W., Wroclawski J., “Assured Forwarding PHB Group”, RFC 2597. June 1999.

[8] Braden B., Clark D., Crowcroft J., Davie B., Deering S., Estrin D., Floyd S., Jacobson V., Minshall G., Partridge C., Peterson L., Ramakrishnan K., Shenker S., Wroclawski J., Zhang L., “Recommendations on Queue Management and Congestion Avoidance in the Internet”, RFC 2309. April 1998.

[9] “Internet Protocol”, RFC 791. September 1981.

[10] Deering S., Hinden R., “Internet Protocol, Version 6 (IPv6) Specification”, RFC 2460. December 1998

[11] Callon R., Doolan P., Feldman N., Fredette A., Swallow G., Viswanathan A., “A Framework for Multiprotocol Label Switching”, draft-ietf-mpls-framework-05.txt., work in progress. September 1999

[12] Le Faucheur, F., Wu L., Davie B., Davari S., Vaananen P., Krishnan R., Cheval P., Heinanen J., “MPLS Support for Differentiated Services”, draft-ietf-mpls-diff-ext-03.txt, work in progress. February 2000.

[13] Awduche D., Berger L., Gan D. Li T., Swallow G., Srinivasan V., “RSVP-TE: Extensions to RSVP for LSP tunnels”, draft-ietf-mpls-rsvp-lsp-tunnel-05.txt, work in progress. February 2000.

[14] Rekhter Y., Rosen E., “Carrying Label Information in BGP-4”, draft-ietf-mpls-bgp4-mpls-04.txt, work in progress. January 2000.

[15] Jamoussi B., “Constraint-Based LSP Setup using LDP”, draft-ietf-mpls-cr-ldp-03.txt, work in progress. September 1999.