

Supporting Differentiated Services in MPLS Networks

Ilias Andrikopoulos and George Pavlou

Centre for Communication Systems Research (CCSR)

University of Surrey

Guildford, Surrey, GU2 5XH, UK

Email: {I.Andrikopoulos, G.Pavlou}@ee.surrey.ac.uk

Abstract: Multi-Protocol Label Switching is a relatively new technology based on the association of labels with routes and the use of labels to forward packets. In other words MPLS integrates the label-swapping paradigm with network-layer routing. Differentiated Services define a model for implementing scalable differentiation of QoS in the Internet. Packets are classified and marked, policed and shaped at the edge of the network in order to receive a particular per-hop forwarding behaviour on nodes along their path. Per-flow state does not need to be maintained in the interior network nodes, thus leading to increased scalability. This obviates the use of complex signalling protocols like RSVP. The inherent characteristics of MPLS make it a very good candidate for providing Differentiated Services. In this paper we describe various approaches which can be used to support differentiated services in MPLS environments.

Keywords: Differentiated Services, Multi-Protocol Label Switching (MPLS), Asynchronous Transfer Mode (ATM), Internet Protocol (IP), Quality of Service (QoS).

1. Introduction

Over the last years a lot of research has been carried out and various standards have been ratified from IETF and ATM Forum addressing the integration of IP and ATM. Example proposed solutions are Classical IP over ATM, Multi-Protocol over ATM (MPOA), LAN Emulation (LANE) and Next Hop Resolution Protocol (NHRP). Additionally, various complex signalling protocols, such as P-NNI, have been developed so that ATM networks can be deployed in the wide area.

MPLS has been recently introduced as a new approach for integrating IP with ATM [1]. Also known as IP switching, IP over ATM, or Layer 3 Switching, it tries to provide the best of both IP and ATM worlds: the efficiency and simplicity of IP routing together with the high-speed switching of ATM by integrating the label-

swapping paradigm with network-layer routing. Label-swapping is performed by associating labels with routes and using the label value to forward packets at Layer 2 of the OSI Reference Model (RM), including the procedure of determining the value of any replacement label. All IP routing functionality remains as is, but the forwarding is now performed at the ATM layer by means of switching. The complex ATM signalling protocols are not required and, more specifically, all the ATM protocols above the ATM Adaptation Layer (AAL) are completely removed.

Although still in the “draft” process within the MPLS Working Group in the IETF, a great deal of research work has been done and several proposals have been submitted. Moreover, a current European ACTS project called IthACI (*Internet and the ATM: Experiments and Enhancements for Convergence and Integration*), aims to provide a number of important enhancements to MPLS: multicast, QoS provisioning, IP mobility and resource management – features which will make MPLS a viable technology. It is in the context of this project the research work described in this paper has been undertaken.

Differentiated Services define a model for implementing scalable differentiation in the Internet. Packets are classified and marked, policed and shaped at the edge of the network in order to receive a particular per-hop forwarding behaviour on nodes along their path. Per-flow state does not need to be maintained in the interior network nodes, which leads to increased scalability.

By closely examining the various characteristics of MPLS, one can see that it is a very good candidate for providing differentiated services. Traffic classification, its ability to reserve Class of Service (CoS) through its lightweight signalling protocol LDP (Label Distribution Protocol) and the label aggregation feature are some of its useful properties.

This paper attempts to show how Differentiated Services can be supported in MPLS networks. Section 2 briefly

presents the main features of the Differentiated Services model and its basic architecture as defined in the current Internet Drafts. Section 3 gives a short introduction to the MPLS architecture and lists some of its main characteristics. In section 4, various approaches for supporting Differentiated Services in MPLS networks are described and a solution is proposed and elaborated. An example is also used to explain analytically the proposed architecture. Finally, our conclusions and a summary are presented in section 5.

2. Differentiated Services

Differentiated services, as proposed by the IETF Differentiated Services Working Group, allow IP traffic to be classified into a finite number of service classes that receive different router treatment. For example, traffic belonging to a higher priority and/or delay service class receives some form of preferential treatment over traffic classified into a lower service class. Differentiated services do not attempt to give explicit end-to-end guarantees. Instead, in congested network elements, traffic with a higher class of priority has a higher probability of getting through, or in case of delay priority, is scheduled for transmission before traffic that is less delay-sensitive [2].

The information required to perform actual differentiation in the network elements is carried in the Type of Service (TOS) field of the IPv4 packet headers or the Traffic Class field of the IPv6 packet headers, referred to as the DS Field or Codepoint (DSCP) [3]. Thus, since the information required by the buffer management and scheduling mechanisms is carried within the packet, differentiated services do not require signalling protocols to control the mechanisms that are used to select different treatment for the individual packets. Consequently, the amount of state information, which is required to be maintained per node, is proportional to the number of service classes and not proportional to the number of application flows.

At each differentiated services user/provider boundary, the service provided is defined by means of a Service Level Agreement (SLA). The SLA is a contract, established either statically or dynamically, that specifies the overall performance and features which can be expected by a customer. Because differentiated services are for unidirectional traffic only, each direction must be considered separately. The subset of the SLA which provides the technical specification of the service is referred to as the Service Level Specification (SLS).

A profound subset of the SLS is the Traffic Conditioning Specification (TCS) which specifies detailed

service parameters for each service level. These service parameters include service performance parameters (e.g. throughput, latency, drop probability) and traffic profiles corresponding to the requested service. Furthermore, the TCS may define the marking and shaping functions to be provided.

2.1 Fundamental Functional Elements of the Differentiated Services Architecture

The Differentiated Services architecture is composed of a number of functional elements, namely packet classifiers, traffic conditioners and per-hop forwarding behaviours (PHB) [4]. According to the basic differentiated services architecture definition, these elements are normally placed in ingress and egress boundary nodes of a differentiated services domain and in interior DS-compliant nodes. However, it is not necessary for all the elements to be present in all the DS-compliant nodes, something that strictly depends on the functionality that is required at each node [5]. In the following paragraphs a short description for each of the elements is given and the various components that comprise them are briefly presented.

Packet Classifiers

Packet classification is a significant function which is normally required at the edge of the differentiated services network. Its goal is to provide identification of the packets belonging to a traffic stream that may receive differentiated services. Classification is done with packet classifiers, which select packets based on the content of packet headers according to well-defined rules determined by the Traffic Conditioning Agreement.

Two types of classifiers are currently defined: the Behaviour Aggregate (BA) classifier, which selects packets based on the DS Codepoint only, and the Multi-Field (MF) classifier, which performs the selection based on the combination of one or more header fields.

Traffic Conditioners

Traffic conditioners form the most vital part of a differentiated services network. Their goal is to apply conditioning functions on the previously classified packets according to a predefined TCS. A traffic conditioner consists of one or more of the following components:

- **Meter**
A device which measures the temporal properties of a traffic stream selected by a classifier.
- **Marker**
A device that sets the DS Codepoint in a packet based on well defined rules.

- **Shaper**

A device that delays packets within a traffic stream to cause the stream to conform to some defined traffic profile.

- **Dropper/Policer**

A device that discards packets based on specified rules (e.g. when the traffic stream does not conform to its TCS).

A typical arrangement of the above mentioned components is illustrated in Figure 1.

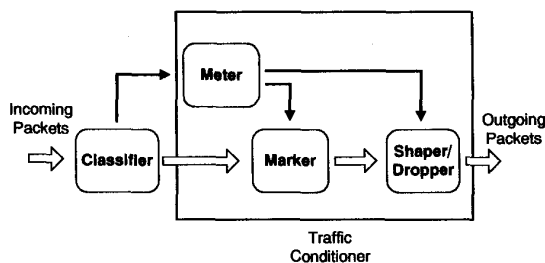


Figure 1 Typical arrangement of a Packet Classifier and a Traffic Conditioner [4].

Per-Hop Forwarding Behaviours (PHB)

A PHB is a description of the externally observable forwarding behaviour of a differentiated services node, applied to a collection of packets with the same DS Codepoint that are crossing a link in a particular direction (called differentiated services behaviour aggregate). Each service class is associated with a PHB. PHBs are defined in terms of behaviour characteristics relevant to service provisioning policies, and not in terms of particular implementations. PHBs may also be specified in terms of their resource priority relative to other PHBs, or in terms of their relative observable traffic characteristics. These PHBs are normally specified as group PHBs and are implemented by means of buffer management and packet scheduling mechanisms.

To preserve partial backwards compatibility with known current uses of the IP Precedence field without sacrificing future flexibility, minimum requirements on a set of PHBs that are compatible with most of the deployed forwarding treatments selected by the IP Precedence field have been defined. In this context, the set of codepoints that are mapped to PHBs meeting these minimum requirements are known as Class Selector Codepoints. The minimum requirements for PHBs that these codepoints may map to are called the Class Selector PHB Requirements. PHBs selected by a Class Selector Codepoint should give

packets a probability of timely forwarding that is not lower than that given to packets marked with a Class Selector codepoint of lower relative order, i.e. smaller numerical value, under reasonable operating conditions and traffic loads [3].

Currently there are three proposed PHBs which are briefly described below.

The Default (DE) PHB is the common, best-effort forwarding available in today's Internet. IP packets marked for this service are sent into a network without adhering to any particular rules and the network will deliver as many of these packets as possible and as soon as possible but without any guarantees.

The Expedited Forwarding (EF) PHB is a high priority behaviour typically used for network control traffic such as routing updates. The EF PHB is defined as a forwarding treatment for a particular differentiated services aggregate where the departure rate of the aggregate's packets from any DS-compliant node must equal or exceed a configurable rate. The EF traffic should be allocated this rate independently of the intensity of any other traffic attempting to transit the node [6].

Finally, the Assured Forwarding (AF) PHB is a means for a provider differentiated services domain to offer different levels of forwarding assurances for IP packets received from a customer differentiated services domain. Four AF classes are defined, where each AF class in each differentiated services node is allocated a certain amount of forwarding resources, e.g. buffer space and bandwidth. Within each AF class, IP packets are marked with one of three possible drop precedence values. In case of congestion, the drop precedence of a packet determines the relative importance of the packet within the AF class [7].

According to the basic architecture assumptions, traffic classifiers and conditioners can be located within DS-compliant nodes at the ingress and egress boundary of a differentiated services domain, although they can also be found in nodes within the interior of a differentiated services domain, or within a non-DS-compliant domain since this is not precluded. However, the exact location of the various components mainly depends on policy and management issues as specified by the network provider.

Typically, end-users/customers will mark their packets to indicate the service they would like to receive. Then, the user traffic entering a differentiated services domain will be conditioned at the ingress node according to the predetermined SLS. Moreover, packets going from one domain to another may need to be re-marked, according to the SLS established between the adjacent domains.

3. Multi-Protocol Label Switching

MPLS is a technology that integrates the label-swapping paradigm with network-layer routing. Although the main focus of MPLS is IP-over-ATM networks, it is not restricted to these technologies. Its goal is to be multi-protocol at both Layer 2 (e.g. ATM, Frame Relay) and Layer 3 (e.g. IP, IPX) of the OSI RM.

Label Switching Routers (LSRs) use link-level forwarding to provide a simple and fast packet-forwarding capability. Label swapping is accomplished by associating fixed-length labels with routes and using the label value to forward packets, including the procedure of determining the value of any replacement label. Depending on the Layer 2 and Layer 3 technologies involved, different label encoding schemes can be used [8]. These are illustrated in Figure 2.

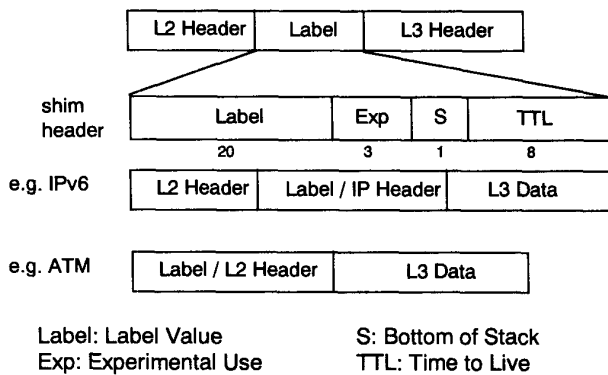


Figure 2 Three different label encoding schemes.

When unlabelled packets need to traverse the same path between an ingress and an egress LSR (packets from an aggregate of one or more flows are said to belong to a stream) belonging to the same MPLS domain, a Label Switched Path (LSP) – a LSP is similar to a unidirectional ATM Virtual Circuit (VC) – needs to be set-up. This will allow the packets to be forwarded from one MPLS node to another just by using the assigned label as an index to a forwarding table. The LSP set-up can be traffic, request, or topology-driven [1]. In the traffic-driven scheme the label assignment is triggered by the arrival of data at an LSR, whereas with the request-driven scheme the label is assigned in response to normal processing of request based control traffic. In the case of a topology-driven scheme the labels are pre-assigned according to existing routing protocol information.

The packets are first classified at the ingress node. Then a mapping between IP packets and a LSP, must take place. This is done by providing a Forwarding Equivalence Class (FEC) specification for each LSP. A FEC is specified as a set of one or more FEC elements, where each FEC element identifies a set of IP packets which may be mapped to the corresponding LSP. Currently, two types of FEC elements exist: the IP address prefix and the host address. In the former, the IP address is said to match the IP address prefix if and only if this address begins with this prefix. In the latter, there must be an exact match between the two addresses.

In the MPLS domain, in order for a LSP to be set-up, labels must be negotiated, distributed, and their semantics defined through a protocol, namely the Label Distribution Protocol (LDP) [9]. LDP is the signalling protocol through which one LSR informs its peers of the label/FEC bindings it has made. An LSR may use a discovery mechanism to discover potential LDP peers. This is done by sending Hello Messages on the MPLS-interface using UDP/IP (User Datagram Protocol / Internet Protocol). Moreover, LDP sessions between LSR peers are established on top of TCP/IP (Transmission Control Protocol / Internet Protocol) -based reliable connections. LDP messages are exchanged through LDP Protocol Data Units (PDUs). Each LDP PDU can carry at least one LDP message. It consists of an LDP header which is followed by one or more LDP messages. The information carried by LDP messages is encoded by using the TLV (Type-Length-Value) scheme. LDP messages are classified under four main categories: discovery, session, advertisement and notification messages.

As the labelled packets are transmitted downstream along the LSP, each LSR examines the label and forwards the packets downstream to the next hop according to its locally significant Next Hop Label Forwarding Entry.




Figure 3 A Multi-Protocol Label Switching network connected to two stub networks on either edge comprising two ingress, two core and two egress Label Switching Routers.

According to Rosen *et al.*, three conceptual information bases are needed to hold MPLS-related information [10]:

- *Next Hop Label Forwarding Entry* (NHLFE). The NHLFE is used when forwarding a labelled packet. It contains the outgoing interface (next hop), the data link encapsulation used for the transmitted packets, the outgoing label and the operation (add, replace, or remove) to perform on the label stack.
- *Incoming Label Map* (ILM). The ILM is a mapping from incoming labels to NHLFEs. It is used when forwarding packets that arrive as labelled packets.
- *FEC-to-NHLFE Map* (FTN). The FTN is a mapping from FECs to NHLFEs. It is used when forwarding packets that arrive unlabeled, but which are to be labelled before forwarding.

In the next section we will be dealing with possible ways for providing support of differentiated services in MPLS networks. These will be further clarified by using an example to describe the operation of the proposed architecture.

4. Differentiated Services and MPLS

As it has already been mentioned in section 2, in order to support differentiated services in a network environment, three fundamental functional elements must be present: packet classifiers, traffic conditioners and per-hop behaviours. We have already discussed how and where these elements should be placed in order for the network to be capable of providing differentiated services. The question that arises is how these components will be efficiently utilised in an MPLS network so that differentiated services are supported.

The support of differentiated services in MPLS environments requires either signalling support for the association of the desired category with the label, or each packet belonging to a stream needs to carry the information of the desired service category (behaviour aggregate).

In this paper we deal with ATM LSRs and hence the packets of a labelled IP stream are actually transported by ATM cells. This poses the question of whether certain peculiarities of ATM should be taken into account or whether a generic approach, independent of the link layer technology, should be followed. If it had not been ATM at Layer 2, it would be possible to include a “shim” header in the packets as mentioned earlier in this paper. However, with ATM, a “shim” header cannot be used because this would involve doing segmentation and re-assembly at

each ATM-LSR in order to read the DSCP field which is against the ATM switching “philosophy”. Hence, the DSCP in the IP header is not accessible by the ATM hardware responsible for the forwarding. Therefore, two alternative solutions may be considered. Either to have some part of the ATM cell header mapped to the DSCP, or to use LDP.

In the first approach, the most likely solution is to use the VPI (Virtual Path Identifier) and part of the VCI (Virtual Channel Identifier) of the ATM cell header as the label, and the remaining eight least significant bits of the VCI be used to map the DSCP [11]. Then all that is needed is the existence of a functional component in the interior DS-compliant ATM LSRs to perform the appropriate traffic management mechanisms on the cells by interpreting the DSCP correctly, with respect to the PHB.

In the second approach, which is more likely for future deployment, the DSCP is mapped to an LSP at the ingress of the MPLS domain. This means that for each DSCP value/PHB a separate LSP will be established for the same egress LSR. So, if there are n Classes and m egress LSRs, $n \cdot m$ LSPs need to be set-up, n labels for each of the m FECs. The packets belonging to streams with the same DSCP and FEC will be forwarded on the same LSP. In other words, the label is regarded as the behaviour aggregate selector.

Furthermore, two LSPs are allowed to be merged into one LSP only if the packets they carry belong to the same Behaviour Aggregate or, even better, if they have the same DSCP. The decision for the merge will be taken at the merging LSR based upon the DSCP entry it has in its modified NHLFE table. Given that the two DSCP values are identical and provided that the necessary resources are available for the rest of the common LSP, the two LSPs can be merged. To check whether there are available resources or not is the role of an admission control module resident in each LSR. A request message needs to be sent to all following hops to check for the necessary bandwidth. If this can be eventually granted, then the merging process may proceed.

Additionally, there must be an MPLS-to-ATM mapping element in every MPLS DS-compliant node which will perform the mapping between the Behaviour Aggregate and the ATM traffic class and traffic parameters.

An issue that would need more discussion is what happens when the MPLS network is topology-driven. Should there be $n \cdot m$ already established LSPs thus forming a kind of overlay network on top of the physical network, or should the LSPs be set-up on demand, which conserves resources in case some of the standard service

classes – and hence the corresponding DSCP values – are rarely used? Evidently, having all LSPs in place is an advantage from the perspective of minimising the LSP set-up delay. Another problem that emerges is the level of aggregation of “microflows” with the same differentiated services behaviour aggregate that can be admitted in such a DS-capable MPLS network. Are the bandwidth reservations per node going to be static or dynamic? If the bandwidth is dynamically allocated, then how will the resources be efficiently partitioned? These are clearly interesting research topics that lie in the areas of resource management and network dimensioning and planning and which are outside the scope of this paper. Here, we make the simple assumption that only best-effort LSPs are initially established and that new LSPs corresponding to specific Behaviour Aggregates need to be set-up.

In the next section, we discuss the modifications and extensions required to be carried out to MPLS.

4.1 Modifications and Extensions to MPLS

In order for MPLS to be able to support differentiated services, a number of modifications/extensions are needed to the LDP protocol and to MPLS in general. These are described below.

First of all, since the MPLS network is considered to be DS-capable, all the functional elements of the differentiated services model must exist and be situated at the same place where they would be in a non-MPLS DS-capable network. The LSRs participating in the MPLS DS-capable network must therefore be DS-compliant. The appropriate PHBs, associated with the various service classes, must also be present in the core DS-compliant LSRs. Given that Layer 2 is ATM, a generic mapping to the corresponding ATM traffic class and parameters is needed. Hence, a mapping element located in the interior nodes will perform the mapping from the currently defined EF, BE and AF classes to ATM. For other types of link layer protocols, suitable mapping elements must exist.

Furthermore, an extension to the Next Hop Label Forwarding Entry (NHLFE) is needed. As stated earlier, in its current form, the NHLFE contains information concerned with forwarding labelled packets and particularly the packet’s next hop (outgoing interface), the data link encapsulation to use when transmitting the packet, the outgoing label and the operation (add, replace, or remove) to perform on the label stack. Moreover, the FTN deals with the forwarding of unlabelled packets. It is therefore necessary to add the DSCP parameter in both the NHLFE and the FTN tables. An example of this mapping table without and with the proposed extension is shown in Figure 4 and Figure 5 respectively.

FEC Element	Port _{in}	VPI _{in}	VCI _{in}	Port _{out}	VPI _{out}	VCI _{out}
a.b.c.d	1	83	95	3	134	198
x.y.w.z	2	45	68	4	129	157

Figure 4 Example of a label table of an FTN without extension for MPLS support (a.b.c.d and x.y.w.z correspond to IP addresses).

FEC Element	Port _{in}	VPI _{in}	VCI _{in}	Port _{out}	VPI _{out}	VCI _{out}	DSCP
a.b.c.d	1	83	95	3	134	198	100101
x.y.w.z	2	45	68	4	129	157	110101

Figure 5 Example of a modified label table of an FTN with extension for MPLS support (a.b.c.d and x.y.w.z correspond to IP addresses).

The next important extension is the addition of the appropriate messages to LDP to make it DS-compliant. There are two basic requirements which need to be fulfilled for this to happen:

- Downstream-on-demand label allocation.
- Addition of the BA attributes in label binding messages.

The first requirement is obvious. In order to set-up end-to-end LSPs with the appropriate differential QoS, we need to ensure that all LSRs belonging to the same LSP perform the label binding in an ordered manner. This can be done by using downstream-on-demand label allocation. The example that follows in the next section shows how this happens.

The way in which the second requirement will be implemented depends on how the differentiated services QoS will be utilised. We propose that the differentiated services QoS is mapped directly to the LDP CoS TLV. The PHB-to-ATM mapper will then be responsible for calculating the necessary QoS parameters (e.g. bandwidth allocation).

Finally, a controller is required to manage and control the ATM switch which forms part of the ATM LSR. Functions such as VC establishment and release, dynamic QoS negotiation, request of switch statistics and configuration information, etc., need to be supported. For this purpose, some kind of general purpose management protocol must be used. An example of such a protocol is Ipsilon’s General Switch Management Protocol (GSMP) [12].

A DS-compliant ATM LSR architecture is illustrated in Figure 6. Its functional elements are described below:

- *TCP/UDP/IP*: This is the TCP/IP protocol stack.
- *MPLS Daemon*: The main process of a LSR. It is where the core of the MPLS protocol is actually located.
- *DS-compatible LDP Daemon*: An LDP daemon process, running on top of TCP/UDP/IP, and which supports the extensions mentioned above. It is used to exchange LDP PDUs with peer LDPs. It also interfaces to the DiffServ module and the MPLS daemon.
- *Admission Control*: It is used to find out whether available resources are sufficient to supply the requested QoS.
- *Routing Daemon*: This is the traditional routing protocol daemon (e.g. OSPF, BGP) running on IP routers.
- *DiffServ Module*: It is responsible for identifying the DSCP at the ingress LSR in order to associate it with the appropriate label. Also, responsible for mapping the PHBs to ATM QoS parameters.
- *Flow MIB*: A database for maintaining flow related information, such as per-flow traffic statistics and path information for aggregated flows. This information is needed for resource management.
- *Flow MIB Controller*: It is responsible for monitoring the LSR and its flows. It collects statistics which are useful for evaluating the local resources.
- *GSMP Interface*: The GSMP protocol is required by the switch controller to control the ATM switch.

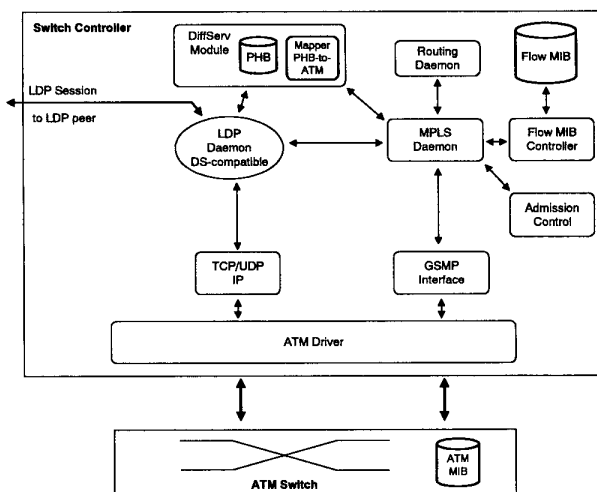


Figure 6 A DS-compliant LSR architecture.

In the following section, we give an example of the operation of a DS-capable MPLS network.

4.2 An Example

We will now explain the LSP set-up procedures for both a non-DS-capable and a DS-capable MPLS network by presenting a detailed example.

We will begin with the description of the default operation in an MPLS network which does not have any differentiated services capabilities. Let's assume that IP traffic belonging to a particular flow and originating from some user at a stub network attached to LSR1, is arriving at LSR1 of the MPLS network which is illustrated in Figure 7. This configuration consists of four edge ATM-LSRs, two ingress and two egress, as well as two core ATM-LSRs and supports topology-driven label assignment and ordered LSP control.




Figure 7 Multi-Protocol Label Switching network with two ingress, two core and two egress Label Switching Routers.

Since the network uses topology-driven label assignment, end-to-end shortcut connections or LSPs from the ingress ATM-LSRs, LSR1 and LSR2, to the egress ATM-LSRs, LSR5 and LSR6, are already in place. The label bindings for this paths will have already been performed through the use of LDP. It should be mentioned, however, that by default the established LSPs are best-effort connections, which in Layer 2, i.e. ATM, context is translated to ATM UBR VCs.

Each IP packet belonging to the same stream is mapped to a corresponding Forwarding Equivalence Class (FEC) when it arrives at LSR1. This FEC has already been assigned a locally significant fixed label. The IP packets are then forwarded to their next hop with the assigned label. At subsequent hops the label is swapped with a new one and the IP packets are forwarded until the egress ATM-LSR where the label is stripped off and then forwarded to the attached stub network.

We will now consider the case where the MPLS network, shown in Figure 7, is DS-capable, hence all required functional elements for providing differentiated

services that have been presented in the previous section, are included in the MPLS nodes. It is assumed that LSPs supporting the various QoS are not set-up in advance and we would like the LSP = (LSR1, LSR3, LSR4, LSR5) to be set-up with a particular QoS.

The first thing to be done is to reserve the necessary bandwidth to accommodate the stream that will be admitted and also allocate the associated labels. The exchange of LDP messages is shown in Figure 8. In case one of the LSRs on the followed path has no adequate resources, it will send a message back to its preceding LSR indicating unavailability of resources. Hence, the LSP path will not be completed.

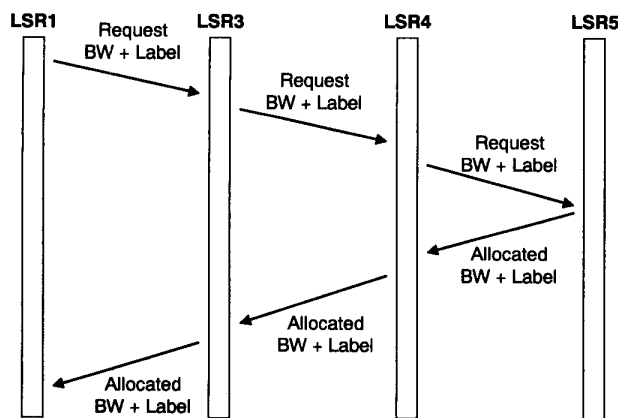


Figure 8 LDP message exchange for requesting and confirming label and bandwidth allocation.

IP packets belonging to a particular traffic stream arrive at LSR1, having already been marked at the source end-host or egress router of the originating network to indicate the level of service they expect. At LSR1, the classification and traffic conditioning functions on the specified traffic are performed by the service provider managing the core network according to a predetermined TCS. Additionally, the network is assumed to have already been provisioned to accept the arriving traffic by statically allocating the necessary resources. The classified IP packets are then checked for their destination IP address and DSCP. These are compared to the entries of the FEC and NHLFE tables. An established LSP which is associated to a FEC element and satisfies the routing and QoS requirements of the stream is found and the corresponding label bound to this LSP is assigned to the IP packets. The rest of the procedure is the same as the one already described earlier in the paper.

5. Conclusions

In this paper, we initially provided an introduction to differentiated services and Multi-protocol Label Switching. We then described the procedures which should be followed and presented the various functional elements which are required in order for Differentiated Services to be supported by MPLS networks. A number of modifications were proposed to be done to the MPLS architecture and its associated signalling protocol, LDP. An example demonstrating how a DS-capable MPLS network works was also given.

We showed how MPLS together with Differentiated Services can be easily combined to form a simple and efficient Internet model capable of providing applications with differential QoS. The need for complex IP and ATM signalling protocols like RSVP and P-NNI respectively is eliminated. No per-flow state information is required leading to increased scalability. A lightweight signalling protocol like LDP with the appropriate extensions along with the ATM traffic management mechanisms, which are already there and implemented in hardware in the ATM switches, provide all the necessary functionality and flexibility required by large networks in a simple manner and without sacrificing precious resources.

Acknowledgements

The opinions presented in this paper are personal opinions and do not represent the official view of the IthACI Consortium. However, we would like to thank our colleagues in the IthACI project for their contribution in formulating some of the views presented here and the European Commission for funding our research activities.

References

- [1] R. Callon, et al., "A Framework for Multiprotocol Label Switching", IETF Internet Draft, work in progress, MPLS WG, November 1997.
- [2] Y. Bernet, et al., "A Framework for Differentiated Services", IETF Internet Draft, work in progress, DiffServ WG, February 1999.
- [3] K. Nichols, et al., "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers", RFC2474, standards track, December 1998.
- [4] S. Blake, et al., "An Architecture for Differentiated Services", RFC2475, informational, December 1998.
- [5] Y. Bernet, et al., "Requirements of DS Boundary Routers", Internet IETF Draft, work in progress, DiffServ WG, February 1999.

- [6] V. Jacobson, et al., "An Expedited Forwarding PHB", IETF Internet Draft, work in progress, DiffServ WG, February 1999.
- [7] J. Heinanen, et al., "Assured Forwarding PHB Group", IETF Internet Draft, work in progress, DiffServ WG, February 1999.
- [8] E. Rosen, et al., "MPLS Label Stack Encoding", IETF Internet Draft, work in progress, MPLS WG, September 1998.
- [9] L. Andersson, et al., "LDP Specification", IETF Internet Draft, work in progress, MPLS WG, January 1999.
- [10] E. Rosen, et al., "Multiprotocol Label Switching Architecture", IETF Internet Draft, work in progress, MPLS WG, February 1999.
- [11] B. Davie, et al., "MPLS using ATM VC Switching", IETF Internet Draft, work in progress, MPLS WG, November 1998.
- [12] P. Newman, et al., "Ipsilon's General Switch Management Protocol Specification Version 2.0", RFC2297, informational, March 1998.