Francois Le Faucheur
Liwen Wu
Bruce Davie
Cisco Systems

Shahram Davari
PMC-Sierra Inc.

Pasi Vaananen
Nokia

Ram Krishnan
Nexabit Networks

Pierrick Cheval
Alcatel

Juha Heinanen
Telia Finland

IETF Internet Draft
Expires: October, 2000
Document: draft-ietf-mpls-diff-ext-03.txt February, 2000

MPLS Support of Differentiated Services

Status of this Memo

Abstract

This document defines a flexible solution for support of Differentiated Services (Diff-Serv) over Multi-Protocol Label Switching (MPLS) networks.

This solution allows the MPLS network administrator to select how Diff-Serv Behavior Aggregates (BAs) are mapped onto Label Switched

Paths so that he/she can best match the Diff-Serv, Traffic Engineering and Fast Restoration objectives within his/her particular network. For instance, this solution allows the network administrator to decide whether different sets of BAs are to be mapped onto the same LSP or mapped onto separate LSPs.

This solution relies on combined use of two types of LSPs:

- LSPs which can transport multiple Ordered Aggregates, so that the EXP field of the MPLS Shim Header conveys to the LSR the PHB to be applied to the packet (covering both information about the packet's scheduling treatment and its drop precedence).
- LSPs which only transport a single Ordered Aggregate, so that the packet's scheduling treatment is inferred by the LSR exclusively from the packet's label value while the packet's drop precedence is conveyed in the EXP field of the MPLS Shim Header or in the encapsulating link layer specific selective drop mechanism (ATM, Frame Relay, 802.1).

1. Introduction

In an MPLS domain [MPLS_ARCH], when a stream of data traverses a common path, a Label Switched Path (LSP) can be established using MPLS signaling protocols. At the ingress Label Switch Router (LSR), each packet is assigned a label and is transmitted downstream. At each LSR along the LSP, the label is used to forward the packet to the next hop.

In a Differentiated Service (Diff-Serv) domain [DIFF_ARCH] all the IP packets crossing a link and requiring the same Diff-Serv behavior are said to constitute a Behavior Aggregate (BA). At the ingress node of the Diff-Serv domain the packets are classified and marked with a Diff-Serv Code Point (DSCP) which corresponds to their Behavior Aggregate. At each transit node, the DSCP is used to select the Per Hop Behavior (PHB) that determines the scheduling treatment and, in some cases, drop probability for each packet.

This document specifies a solution for supporting the Diff-Serv Behavior Aggregates whose corresponding PHBs are currently defined (in [DIFF_HEADER], [DIFF_AF], [DIFF_EF]) over an MPLS network. This solution also offers flexibility for easy support of PHBs that may be defined in the future.

As mentioned in [DIFF_HEADER], "Service providers are not required to use the same node mechanisms or configurations to enable service differentiation within their networks, and are free to configure the node parameters in whatever way that is appropriate for their service offerings and traffic engineering objectives". Thus, the solution defined in this document gives Service Providers flexibility in selecting how Diff-Serv classes of service are Routed

or Traffic Engineered within their domain (eg. separate classes of services supported via separate LSPs and Routed separately, all classes of service supported on the same LSP and Routed together). Similarly, the solution gives Service Providers flexibility in how Diff-Serv classes of service can be protected via MPLS Fast

Restoration (eg. some classes of service supported via LSPs which are protected via MPLS Fast Restoration while some other classes of service are supported via LSPs which are not protected).

Beside, the solution specified in this document achieves label space conservation and reduces the volume of label set-up/tear-down signaling where possible by only resorting to multiple LSPs for a given Forwarding Equivalent Class (FEC) [MPLS_ARCH] when useful or required.

This specification allows support of Differentiated Services for both IPv4 and IPv6 traffic transported over an MPLS network.

This document only describes operations for unicast. Multicast support is for future study

1.1 Ordered Aggregate (OA) and PHB Scheduling Class (PSC)

The Diff-Serv model defines [DIFF_NEW] the set of Behavior Aggregates which share an ordering constraint to constitute an "Ordered Aggregate (OA)". It also defines the set of one or more PHBs that are applied to this set of Behavior Aggregates to constitute a "PHB Scheduling Class (PSC)".

1.2 EXP-Inferred-PSC LSPs (E-LSP)

A single LSP can be used to support up to eight BAs of a given FEC, regardless of how many OAs these BAs span. With such LSPs, the EXP field of the MPLS Shim Header [MPLS_ENCAPS] is used by the LSR to determine the PHB to be applied to the packet. This includes both the PSC and the drop preference.

We refer to such LSPs as "EXP-inferred-PSC LSPs" (E-LSP), since the PSC of a packet transported on this LSP depends on the EXP field value for that packet.

The mapping from EXP field to PHB (ie to PSC and drop precedence) for a given such LSP, is either explicitly signaled at label set-up or relying on a pre-configured mapping.

Detailed operations of E-LSPs are specified in section 3 below.

1.3 Label-Only-Inferred-PSC LSPs (L-LSP)

A separate LSP can be established for a single <FEC, OA> pair.

With such LSPs, the PSC is explicitly signaled at label establishment time so that, after label establishment, the LSR can infer exclusively from the label value the PSC to be applied to a labeled packet. When the Shim Header is used, the Drop Precedence to be applied by the LSR to the labeled packet, is conveyed inside the labeled packet MPLS Shim Header using the EXP field [MPLS_ENCAPS]. When the Shim Header is not used (eg. MPLS Over ATM), the Drop Precedence to be applied by the LSR to the labeled packet is

conveyed inside the link layer header encapsulation using link layer specific drop precedence fields (eg. ATM Cell Loss Priority).

We refer to such LSPs as "Label-Only-Inferred-PSC LSPs" (L-LSP) since the PSC can be fully inferred from the label without any other information (eg. regardless of the EXP field value). Detailed operations of L-LSPs are specified in section 4 below.

1.4 Overall Operations

For a given FEC, and unless media specific restrictions apply as identified in the sections 7, 8, 9 and 10 below, this specification allows any one of the following combinations within an MPLS Diff-Serv domain:

- zero or any number of E-LSPs, and
- zero or any number of L-LSPs.

The network administrator selects the actual combination of LSPs from the set of allowed combinations and selects how the Behavior Aggregates are actually transported over this combination of LSPs, in order to best match his/her environment and objectives in terms of Diff-Serv support, Traffic Engineering and Fast Restoration. Criteria for selecting such a combination are outside the scope of this specification; However in order to respect ordering constraints, all packets of a given microflow, possibly spanning multiple BAs of a given Ordered Aggregate, MUST be transported over the same LSP. Conversely, each LSP MUST be capable of supporting all the (active) PHBs of a given PSC.

Examples of deployment scenarios are provided for information in APPENPIX A.

1.5 Relationship between Label and FEC

[MPLS_ARCH] states in section ‘2.1. Overview’ that:
‘Some routers analyze a packet’s network layer header not merely to choose the packet’s next hop, but also to determine a packet’s "precedence" or "class of service". They may then apply different discard thresholds or scheduling disciplines to different packets. MPLS allows (but does not require) the precedence or class of service to be fully or partially inferred from the label. In this case, one may say that the label represents the combination of a FEC and a precedence or class of service.’

In line with this, we observe that:

- With E-LSPs, the label represents the combination of a FEC and the set of Behavior Aggregates (BAs) transported over the E-LSP). Where all the supported BAs are transported over an E-LSP, the label then represents the complete FEC.
- With L-LSPs, the label represents the combination of a FEC and an Ordered Aggregate (OA).

2. Label Forwarding Model for Diff-Serv LSRs

Since different Ordered Aggregates of a given FEC may be transported over different LSPs, the label swapping decision of a Diff-Serv LSR clearly depends on the forwarded packet's Behavior Aggregate. Also, since the IP DS field of a forwarded packet may not be directly visible to an LSR, the way to determine the PHB to be applied to a received packet and to encode the PHB into a transmitted packet is different to a non-MPLS Diff-Serv Router.

In order to describe Label Forwarding by Diff-Serv LSRs, we model the LSR Diff-Serv label switching behavior as comprising four stages:

- Incoming PHB Determination (A)
- Optional Outgoing PHB Determination via Local Policy and Traffic Conditioning (B)
- Label Swapping (C)
- Encoding of Diff-Serv information into Encapsulation Layer (EXP,CLP,DE,User_Priority) (D)

Obviously, to enforce the Diff-Serv service differentiation the LSR MUST also apply the forwarding treatment corresponding to the Outgoing PHB.

This model is illustrated below:

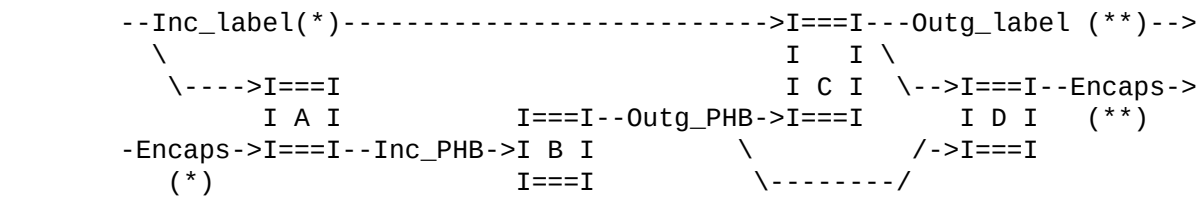

'Encaps' designates the Diff-Serv related information encoded in the MPLS Encapsulation layer (eg EXP field, ATM CLP, Frame Relay DE, 802.1 User_Priority)

(*) when the LSR performs label imposition, the incoming packet is received unlabelled.

(**) when the LSR performs label disposition, the outgoing packet is transmitted unlabelled.

This model is presented here to illustrate operations of Diff-Serv LSRs and does not constrain actual implementation.

2.1 Incoming PHB Determination

This stage determines which Behavior Aggregate the received packet belongs to.

2.1.1 Incoming PHB Determination for received labelled packets

This specification defines one default method for this determination

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.