Network Working Group
Request for Comments: 3214
Category: Standards Track

J. Ash
AT&T
Y. Lee
Ceterus Networks
P. Ashwood-Smith
B. Jamoussi
D. Fedyk
D. Skalecki
Nortel Networks
L. Li
SS8 Networks
January 2002

LSP Modification Using CR-LDP

Status of this Memo

Copyright Notice

Abstract

This document presents an approach to modify the bandwidth and possibly other parameters of an established CR-LSP (Constraint-based Routed Label Switched Paths) using CR-LDP (Constraint-based Routed Label Distribution Protocol) without service interruption. After a CR-LSP is set up, its bandwidth reservation may need to be changed by the network operator, due to the new requirements for the traffic carried on that CR-LSP. The LSP modification feature can be supported by CR-LDP by use of the _modify_value for the _action indicator flag_ in the LSPID TLV. This feature has application in dynamic network resources management where traffic of different priorities and service classes is involved.

Table of Contents

1. Conventions Used in This Document

L:	LSP (Label Switched Path)
L-id:	LSPID (LSP Identifier)
T:	Traffic Parameters
R:	LSR (Label Switching Router)
FEC:	Forwarding Equivalence Class
NHLFE:	Next Hop Label Forwarding Entry
FTN:	FEC To NHLFE
TLV:	Type Length Value

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [4].

2. Introduction

Consider an LSP L1 that has been established with its set of traffic parameters T0. A certain amount of bandwidth is reserved along the path of L1. Consider then that some changes are required on L1. For example, the bandwidth of L1 needs to be increased to accommodate the increased traffic on L1. Or the SLA associated with L1 needs to be modified because a different service class is desired. The network operator, in these cases, would like to modify the characteristics of L1, for example, to change its traffic parameter set from T0 to T1, without releasing the LSP L1 to interrupt the service. In some other cases, network operators may want to reroute a CR-LSP to a different path for either improved performance or better network resource utilization. In all these cases, LSP modification is required. In section 3 below, a method to modify an active LSP using CR-LDP is

presented. The concept of LSPID in CR-LDP is used to achieve the LSP modification, without releasing the LSP and interrupting the service and, without double booking the bandwidth. In Section 4, an example is described to demonstrate an application of the presented method in dynamically managing network bandwidth requirements without interrupting service. In CR-LDP, an action indicator flag of _modify_ is used in order to explicitly specify the behavior, and allow the existing LSPID to support other networking capabilities in the future. Reference [3], RFC XXXX, specifies the action indicator flag of _modify_ for CR-LDP.

3. LSP Modification Using CR-LDP

3.1 Basic Procedure for Resource Modification

LSP modification can only be allowed when the LSP is already set up and active. That is, modification is not defined nor allowed during the LSP establishment or label release/withdraw phases. Only modification requested by the ingress LSR of the LSP is considered in this document for CR-LSP. The Ingress LSR cannot modify an LSP before a previous modification procedure is completed.

Assume that CR-LSP L1 is set up with LSPID L-id1, which is unique in the MPLS network. The ingress LSR R1 of L1 has in its FTN (FEC To NHLFE) table FEC1 -> Label A mapping where A is the outgoing label for LSP L1. To modify the characteristics of L1, R1 sends a Label Request Message. In the message, the TLVs will have the new requested values, and the LSPID TLV is included which indicates the value of L-id1. The Traffic Parameters TLV, the ER-TLV, the Resource Class (color) TLV and the Preemption TLV can have values different from those in the original Label Request Message, which has been used to set up L1 earlier. Thus, L1 can be changed in its bandwidth request (traffic parameter TLV), its traffic service class (traffic parameter TLV), the route it traverses (ER TLV) and its setup and holding (Preemption TLV) priorities. The ingress LSR R1 now still has the entry in its FTN as FEC1 -> Label A. R1 is waiting to establish another entry for FEC1.

When an LSR Ri along the path of L1 receives the Label Request message, its behavior is the same as that of receiving any Label request message. The only extension is that Ri examines the LSPID carried in the Label Request Message, L-id1, and identifies if it already has L-id1. If Ri does not have L-id1, Ri behaves the same as receiving a new Label Request message. If Ri already has L-id1, Ri takes the newly received Traffic Parameter TLV and computes the new bandwidth required and derives the new service class. Compared with the already reserved bandwidth for L-id1, Ri now reserves only the difference of the bandwidth requirements. This prevents Ri from doing

bandwidth double booking. If a new service class is requested, Ri also prepares to receive the traffic on L1 in just the same way as handling it for a Label Request Message, perhaps using a different type of queue. Ri assigns a new label for the Label Request Message.

When the Label Mapping message is received, two sets of labels exist for the same LSPID. Then the ingress LSR R1 will have two outgoing labels, A and B, associated with the same FEC, where B is the new outgoing label received for LSP L1. The ingress LSR R1 can now activate the new entry in its FTN, FEC1 - > Label B. This means that R1 swaps traffic on L1 to the new label _B_ (_new_ path) for L1. The packets can now be sent with the new label B, with the new set of traffic parameters if any, on a new path, that is, if a new path is requested in the Label Request Message for the modification. All the other LSRs along the path will start to receive the incoming packets with the new label. For the incoming new label, the LSR has already established its mapping to the new outgoing label. Thus, the packets will be sent out with the new outgoing label. The LSRs do not have to implement new procedures to track the new and old characteristics of the LSP.

The ingress LSR R1 then starts to release the original label A for LSP L1. The Label Release Message is sent by R1 towards the down stream LSRs. The Release message carries the LSPID of L-id1 and the Label TLV to indicate which label is to be released. The Release Message is propagated to the egress LSR to release the original labels previously used for L1. Upon receiving the Label Release Message, LSR Ri examines the LSPID, L-id1, and finds out that the L-id1 has still another set of labels (incoming/outgoing) under it. Thus, the old label is released without releasing the resource in use. That is, if the bandwidth has been decreased for L1, the delta bandwidth is released. Otherwise, no bandwidth is released. This modification procedure can not only be applied to modify the traffic parameters and/or service class of an active LSP, but also to reroute an existing LSP (as described in Section 3.2 below), and/or change its setup/holding priority if desired. After the release procedure, the modification of the LSP is completed.

The method described above follows the normal behavior of Label Request / Mapping / Notification / Release / Withdraw procedure of a CR-LDP operated LSR with a specific action taken on an LSPID. If a Label Withdraw Message is used to withdraw a label associated with an LSPID, the Label TLV should be included to specify which label to withdraw. Since the LSPID can also be used for other feature support, an action indication flag of _modify_ assigned to the LSPID would explicitly explain the action/semantics that should be associated with the messaging procedure. The details of this flag are addressed in the CR-LDP document, Reference [3].

3.2 Rerouting LSPs

LSP modification can also be used to reroute an existing LSP. Only modification requested by the ingress LSR of the LSP is considered in this document for CR-LSP. The Ingress LSR cannot modify an LSP before a previous modification procedure is completed.

As in the previous section, consider a CR-LSP L1 with LSPID L-id1. To modify the route of the LSP, the ingress LSR R1 sends a Label Request Message. In the message, the LSPID TLV indicates L-id1 and the Explicit Route TLV is specified with some different hops from the explicit route specified in the original Label Request Message. The action indication flag has the value _modify_.

At this point, the ingress LSR R1 still has an entry in FTN as FEC1 -> Label A. R1 is waiting to establish another entry for FEC1.

When an LSR Ri along the path of L1 receives the Label Request message, its behavior is the same as that of receiving a Label Request Message that modifies some other parameters of the LSP. Ri assigns a new label for the Label Request Message and forwards the message along the explicit route. It does not allocate any more resources except as described in section 3.1.

At another LSR Rj further along the path, the explicit route diverges from the previous route. Rj acts as Ri, but forwards the Label Request message along the new route. From this point onwards the Label Request Message is treated as setting up a new LSP by each LSR until the paths converge at later LSR Rk. The _modify_ value of the action indication flag is ignored.

At Rk and subsequent LSRs, the Label Request Message is handled as at Ri.

On the return path, when the Label Mapping message is received, two sets of labels for the LSPID exist where the new route coincide with the old. Only one set of labels will exist at LSRs where the routes diverge.

When the Label Mapping message is received at the ingress LSR R1 it has two outgoing labels, A and B, associated with the same FEC, where B is the new outgoing label received for LSP L1. R1 can now activate the new entry in the FTN, FEC1 - > Label B and de-activate the old entry FEC1 - > Label A. This means that R1 swaps traffic on L1 to the new label B. The packets are now sent with the new label B, on the new path.

The ingress LSR R1 then starts to release the original label A for LSP L1. The Label Release Message is sent by R1 towards the down stream LSRs following the original route. The Release message carries the LSPID of L-id1 and the Label TLV to indicate which label is to be released. At each LSR the old label is released - no further action is required to change the path of the data packets which are already following the new route programmed by the Label Mapping message.

At some LSRs, where the routes diverged, there is only one label for the LSPID. For example, between Rj and Rk, the Label Release Message will follow the old route. At LSRs between Rj and Rk only the labels from the original route will exist for LSPID L-id1. At these LSRs the LSPID TLV does not need to be examined to release the correct label, but it must still be updated and passed on to the next LSR as the Label Release message is propagated. In this way, at Rk where the routes converge, the downstream LSR will know which label to release and can continue to forward the Label Release Message along the old route.

3.3 Priority Handling

When sending a Label Request Message for an active LSP L1 to request changes, the setup priority used in the label Request Message can be different from the one used in the previous Label Request Message, effectively indicating the priority of this _modification_ request. Network operators can use this feature to decide what priority is to be assigned to a modification request, based on their policies/algorithms and other traffic situations in the network. For example, the priority for modification can be determined by the priority of the customer/LSP. If a customer has exceeded the reserved bandwidth of its VPN LSP tunnel by too much, the modification request's priority may be given as a higher value. The Label Request message for the modification of an active LSP can also be sent with a holding priority different from its previous one. This effectively changes the holding priority of the LSP. Upon receiving a Label Request Message that requests a new holding priority, the LSR assigns the new holding priority to the bandwidth. That is, the new holding priority is assigned to both the existing incoming / outgoing labels and the new labels to be established for the LSPID in question. In this way self-bumping is prevented.

3.4 Modification Failure Case Handling

A modification attempt may fail due to insufficient resource or other situations. A Notification message is sent back to the ingress LSR R1 to indicate the failure of Label Request Message that intended to modify the LSP. A retry may be attempted if desired by the network

operator. If the LSP on the original path failed when a modification attempt is in progress, the attempt should be aborted by using the Label Abort Request message as specified in the LDP document [5].

In the event of a modification failure, all modifications to the LSP including the holding priority must be restored to their original values.

4. Application of LSP Bandwidth Modification in Dynamic Resource Management

In this section, we gave an example of dynamic network resource management using the LSP bandwidth modification capability. The details of this example can be found in a previous internet-draft [2]. Assume that customers or services are assigned with given CR-LSPs. These customers/services are assigned with one of three priorities: key, normal or best effort. The network operator does not want to bump any LSPs during an LSP setup, so after these CR-LSPs are set up, their holding priorities are all assigned as the highest value.

The network operator wants to control the resource on the links of the LSRs, so each LSR keeps the usage status of its links. Based on the usage history, each link is assigned a current threshold priority Pi, which means that the link has no bandwidth available for a Label Request with a setup priority lower than Pi. When an LSP's bandwidth needs to be modified, the operator uses a policy-based algorithm to assign a priority for its modification request, say Mp for LSP L2. The ingress LSR then sends a Label Request message with Setup Priority = Mp. If there is sufficient bandwidth on the link for the modification, and the Setup priority in the Label Request Message is higher in priority (Mp numerically smaller) than the Pi threshold of the link, the Label Request Message will be accepted by the LSR. Otherwise, the Label Request message will be rejected with a Notification message which indicates that there are insufficient resources. It should also be noted that when OSPF (or IS-IS) floods the available-link-bandwidth information, the available bandwidth associated with a priority lower than Pi (numerical value bigger) should be interpreted as _0_.

This example based on a priority threshold Pi is implementation specific, and illustrates the flexibility of the modification procedure to prioritize and control network resources. The calculation of Mp can be network and service dependent, and is based on the operator's routing policy. For example, the operator may assign a higher priority (lower Mp value) to L2 bandwidth modification if L2 belongs to a customer or service with _Key_ priority. The operator may also collect the actual usage of each LSP

and assign a lower priority (higher Mp) to L2 bandwidth-increase modification if, for example, in the past week L2 has exceeded its reserved bandwidth by 2 times on the average. In addition, an operator may try to increase the bandwidth of L2 on its existing path unsuccessfully if there is insufficient bandwidth available on L2. In that case, the operator is willing to increase the bandwidth of another LSP, L3, with the same ingress/egress LSRs as L2, in order to increase the overall ingress/egress bandwidth allocation. However, in this case the L3 bandwidth modification is performed with a lower priority (higher Mp value) since L3 is routed on a secondary path, which results in the higher bandwidth allocation priority being given to the LSPs that are on their primary paths [2].

5. Acknowledgments

The authors would like to acknowledge the careful review and comments of Adrian Farrel.

6. Intellectual Property Considerations

The IETF has been notified of intellectual property rights claimed in regard to some or all of the specification contained in this document. For more information consult the online list of claimed rights.

7. Security Considerations

Protection against modification to LSPs by malign agents has to be controlled by the MPLS domain.

8. References

[1] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

[2] Ash, J., "Traffic Engineering & QoS Methods for IP-, ATM-, & TDM-Based Multiservice Networks", Work in Progress.

[3] Jamoussi, B., Editor, Andersson, L., Callon, R., Dantu, R., Wu, L., Doolan, P., Worster, T., Feldman, N., Fredette, A., Girish, M., Gray, E., Heinanen, J., Kilty, T. and A. Malis, "Constraint-based LSP Setup Using LDP", RFC 3212, January 2002.

[4] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[5] Andersson, L., Doolan, P., Feldman, N., Fredette, A. and B. Thomas, "LDP Specification", RFC 3036, January 2001.

[6] Rosen, E., Viswanathan, A. and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[7] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M. and J. McManus, "Requirements for Traffic Engineering Over MPLS", RFC 2702, September 1999.

[8] Ash, J., Girish, M., Gray, E., Jamoussi,B. and G. Wright, "Applicability Statement for CR-LDP", RFC 3213, January 2002.

9. Authors' Addresses

Gerald R. Ash
AT&T
Room MT D5-2A01
200 Laurel Avenue
Middletown, NJ 07748
USA
Phone: 732-420-4578
EMail: gash@att.com

Bilel Jamoussi
Nortel Networks Corp.
600 Tech Park
Billerica, MA 01821
USA
Phone: 978-288-4506
EMail: jamoussi@NortelNetworks.com

Peter Ashwood-Smith
Nortel Networks Corp.
P O Box 3511 Station C
Ottawa, ON K1Y 4H7
Canada
Phone: +1 613 763-4534
EMail: petera@NortelNetworks.com

Darek Skalecki
Nortel Networks Corp.
P O Box 3511 Station C
Ottawa, ON K1Y 4H7
Canada
Phone: +1 613 765-2252
EMail: dareks@nortelnetworks.com

Young Lee
Ceterus Networks
EMail: ylee@ceterusnetworks.com

Li Li
SS8 Networks
495 March Rd., 5th Floor
Kanata, Ontario
K2K 3G1 Canada
Phone: +1 613 592-2100 ext. 3228
EMail: lili@ss8networks.com

Don Fedyk
Nortel Networks Corp.
600 Tech Park
Billerica, MA 01821
USA
Phone: 978-288-3041
EMail: dwfedyk@nortelnetworks.com

10. Full Copyright Statement

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.