Network Working GroupRohit Dube (Xebeo Communications)Internet DraftMichele Costa (Xebeo Communications)Expiration Date: January 2004July 2003

Bi-directional LSPs for classical MPLS

draft-dube-bidirectional-lsp-01.txt

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of <u>Section 10 of RFC2026</u>.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/lid-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

Abstract

This memo proposes extensions to support bi-directional LSPs for classical MPLS. Specifically RSVP-TE is extended with objects similar to those in GMPLS to allow establishment of bi-directional LSPs for MPLS networks.

1. Introduction

Bi-directional LSPs are well known in a GMPLS context. However, no mechanism exists to establish bi-directional LSPs with with classical (non-generalized) MPLS labels. Allowing bi-directional LSPs in MPLS networks is useful for operators from a maintenance point of view because bi-directional LSPs traverse the same path through the network in both directions. This simplifies the operators view of the network as well as path failure modes - since the forward and reverse paths are identical. Further, bi-directional LSPs require less state when compared to two uni-directional LSPs [GMPLS-SIG].

Bi-directional LSPs are especially relevant when MPLS is used to provide a service which requires connectivity in both directions between two LERs. Applications include pseudo-wire emulation [PWE3] of Ethernet, Frame Relay and ATM.

In this memo, we use the term MPLS to mean "classical" MPLS and GMPLS to mean Generalized MPLS. Further, when describing bi-directional LSPs, we will use the term "initiator" to represent the LER that initiates LSP setup and "terminator" to represent the LER at the other end of the LSP.

Expires January 2004

2. Details

In order to support bi-directional LSPs, traffic engineered paths in both the forward and reverse directions need to be calculated and the signaling protocol needs to carry additional labels for the reverse path. For RSVP-TE, extensions are needed to carry MPLS labels in PATH messages, such that the reverse path is setup at the same time as the forward path.

2.1 CSPF extensions

Constrained Shortest-Path-First (CSPF) calculations in MPLS networks are usually in the forward direction from an ingress LER. This is adequate for IP traffic-engineering [TE] due to the asymmetrical nature of the underlying IP traffic. For applications like pseudo-wire transport over MPLS, a symmetric path is often desired.

To support this functionality, the CSPF implementation at an LER needs to be able to calculate a strict and explicit path which meets the TE constraints in both in the forward and reverse directions.

2.1.1 CSPF extensions for asymmetric bandwidth reservation

An application may desire asymmetric bandwidth reservation in the forward and reverse directions. To accommodate this the CSPF implementation should be capable of accepting separate sets of constraints for the forward and reverse directions.

2.2 RSVP-TE extensions

[GMPLS-RSVP-TE] describes bi-directional LSP setup in a GMPLS environment. The primary mechanism used for this support is the UPSTREAM_LABEL object which is sent out as part of the PATH message and contains the label to be used in the reverse direction.

[GMPLS-RSVP-TE] defines the use of UPSTREAM_LABEL objects only with generalized labels (c-type=2). This object needs to be implemented

for MPLS with regular labels (c-type=1).

2.2.1 RSVP-TE extensions for asymmetric bandwidth reservation

An Asymmetric bandwidth reservation request is indicated by the presence of an UPSTREAM_FLOWSPEC object. The initiator includes an UPSTREAM_FLOWSPEC object describing the QoS properties for the reverse direction in the PATH message. The format of this object is the same as the RSVP FLOWSPEC object [<u>RSVP-INTSERV</u>].

An implementation supporting the UPSTREAM_FLOWSPEC object should allocate any necessary resources before sending a PATH message to a node downstream.

Expires January 2004

3. Operation

On receiving an operators input, the initiator LER kicks off a a path computation which satisfies the TE constraints in both the forward and reverse directions. The resulting strict and explicit hop list is handed to RSVP-TE for a bi-directional LSP setup.

The initiator populates a label in the UPSTREAM_LABEL object which it sends as part of the PATH message to the downstream node. The downstream node uses the label in the UPSTREAM_LABEL object for the reverse direction. When sending the UPSTREAM_LABEL object to the next node further downstream, it similarly populates the UPSTREAM_LABEL object with an MPLS label for use in the reverse direction. This process stops when the terminator receives the PATH message.

The terminator then sends back a RESV message (which is no different from that in [RSVP-TE], [RSVP]) to setup the LSP in the forward direction from the terminator. At the end of this transaction, a bi-directional LSP is setup requiring no more messages than that needed for a regular uni-directional LSP.

4. Conclusion

Bi-directional LSPs are a useful construct which simplify operations wherever connectivity in both directions is required. They are efficient as they require the same number of messages as regular uni-directional LSPs for path maintenance. Implementation experience shows that the memory requirement for a bi-directional lsp when compared to two uni-directional LSPs is lower as less state needs to be maintained for one bi-directional LSP when compared to two uni-directional LSPs Extending their applicability from GMPLS to MPLS is straight-forward and the implementation overhead is low. Since CR-LDP [<u>CR-LDP</u>] supports bi-directional LSPs for GMPLS as well [<u>GMPLS-CR-LDP</u>], enhancements similar to those in <u>section 2.2</u> can be made to support bi-directional LSPs via CR-LDP for MPLS.

5. Security Considerations

The procedures and extensions proposed in this document do not raise any new security concerns.

<u>6</u>. IANA Considerations

For label allocation in the reverse direction, there are no additional IANA issues besides those already present in [<u>GMPLS-RSVP-TE</u>]. The class-num value assigned to the UPSTREAM_LABEL object will work as is for any value assigned by IANA [<u>IANA</u>].

For asymmetric bandwidth reservation on the revere path a new RSVP Class-Num is needed for the UPSTREAM_FLOWSPEC object. A value of 38 is suggested.

Expires January 2004

7. Acknowledgments

We would like to thank Giles Heron, Pascal Menezes, Lior Shabtay, Yaron Raz, Vasanthi Thirumalai, Dawn Xie and Kerry Fendick for their comments on this memo.

8. References

[CR-LDP] B. Jamoussi, Ed. et al, "Constraint-Based LSP Setup using LDP", <u>RFC 3212</u>.

[IANA] T. Narten and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", <u>RFC 2434</u>.

[GMPLS-SIG] L. Berger, Ed., et al, "Generalized MPLS - Signaling Functional Description", Internet Draft, <u>draft-ietf-mpls-generalized-signaling-09.txt</u>

[GMPLS-CR-LDP] P. Ashwood-Smith, Ed., et al, "Generalized MPLS Signaling - CR-LDP Extensions", Internet Draft, <u>draft-ietf-mpls-generalized-cr-ldp-06.txt</u>.

[GMPLS-RSVP-TE] L. Berger, Ed., et al, "Generalized MPLS Signaling -RSVP-TE Extensions", Internet Draft, <u>draft-ietf-mpls-generalized-rsvp-te-09.txt</u>.

[PWE3] X. Xiao, et al, "Requirements for Pseudo-Wire Emulation Edge-to-Edge", Internet Draft, <u>draft-ietf-pwe3-requirements-01.txt</u> [RSVP] R. Braden, Ed., et al, "Resource ReSerVation protocol (RSVP) -- version 1 functional specification", <u>RFC2205</u>.

[RSVP-INTSERV] R. Braden, Ed., et al, "The Use of RSVP with IETF Integrated Services", <u>RFC2210</u>.

[RSVP-TE] D. Awduche, et al, "RSVP-TE: Extensions to RSVP for LSP tunnels", <u>RFC 3209</u>.

[TE] D. Awduche, et al, "Requirements for Traffic Engineering Over MPLS", <u>RFC 2702</u>.

9. Author Information

Rohit Dube Xebeo Communications Inc. 1 Cragwood Road, Suite 100 South Plainfield, NJ 07080 e-mail: rohit@xebeo.com

Michele Costa Xebeo Communications Inc. 1 Cragwood Road, Suite 100 South Plainfield, NJ 07080 e-mail: mcosta@xebeo.com