MPLS Working Group
Internet Draft
Intended status: Informational
Expires: May 2012

Dave Allan, Greg Mirsky
Ericsson

November 2011

A framework for the use of SPMEs for shared mesh protection
draft-allan-spme-smp-fmwk-00

Abstract

Shared mesh protection allows a set of diversely routed paths with diverse endpoints to collectively oversubcribe protection resources. Under normal conditions no single failure will result in the capacity of the associated protection resources to be exhausted.

When multiple failures occur such that more than one path in the set of paths utilizing shared protection resources is affected, the necessity arises of pre-empting traffic on the basis of business priority rather than application priority.

This memo describes the use of SPMEs and TC marking as a means of indicating business priority for shared mesh protection.

Requirements Language

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 RFC2119 [1].

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

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/1id-abstracts.txt.

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

This Internet-Draft will expire on August 2nd 2011.

Copyright Notice

Table of Contents

1. Introduction

Shared mesh protection is described in [2]. A common interpretation of the behavior of shared mesh protection emerges from the circuit switched world whereby subtending path selectors and selector coordination functions support path preemption to ensure that the highest priority path needing the protection resources is granted ownership of the shared segment, all others being preempted, and such functionality can be successfully delegated to dataplane OAM.

Ultimately this resolves into a business priority decision vs. an application priority decision in how customer traffic is handled. The packet world is different from the circuit world in that there is no guarantee of convenient alignment of resource requirements between preempting and preempted paths. Nor in a packet environment is there the need to completely preempt all the traffic in a lower priority path simply because a higher priority path lays claim to the resources. Finally it is useful to obviate the requirement for preempting and preemptable traffic to be co-routed.

This memo proposes the use of SPMEs with the pipe model of TC copying as an alternative to the use of path pre-emption, path selectors and selector coordination functions for the purposes of implementing business policy.

1.1. Authors

David Allan, Greg Mirsky

2. Conventions used in this document

2.1. Terminology

MPLS-TP: MPLS Transport Profile

MPLS-TP LSP: Uni-directional or Bidirectional Label Switch Path representing a circuit

SMP: Shared Mesh Protection

SPME: Sub-Path Maintenance Entity

TC: Traffic Class

TTL: Time To Live

3. Overview

Shared mesh protection is described in [2]. A common interpretation of the behavior of shared mesh protection emerges from the circuit switched world. In that interpretation subtending path selectors and selector coordination support path preemption functionality to ensure that the highest priority path needing the protection resources is the one granted ownership of the shared segment; all others being preempted. It also assumes that all paths sharing the protection resources conveniently all need exactly the same size pipe.

In packet transport networks there will frequently not be a convenient 1:1 equivalence of the bandwidth requirements of the set of transport paths sharing protection resources such that a simple pre-emption decision can be made. For example 3 paths: A, B, and C sized "n", "n/2" and "n/2" respectively could have a shared segment size "3n/2" such that simultaneous failures necessitating the activation of any two of the protection paths could be accommodated without path preemption. When one ranks A, B and C with a variety of priorities and considers all failure combinations a rather large matrix of possible required behaviors emerges.

If one pursues this line of thinking to its logical conclusion, and envisions a significant set of paths of diverse sizes and diverse priorities, the policy associated with successful path prioritization and preemption becomes quite complex, and ensuring multiple selectors make timely and of necessity common preemption decisions starts to impose network design constraints or require additional coordination protocols that severely impact the utility of SMP.

Further in a packet network there can be a difference in the bandwidth reserved and the bandwidth actually used at any given instant in time. One consequence is that there is no need to completely preempt all the traffic in a lower priority path simply because a higher priority path lays a preferential claim to the bandwidth.

To obviate these problems, this memo proposes an alternative to how business priority can be implemented for shared mesh protection that obviates the need for path preemption and the limitations such an approach imposes.

3.1. Architectural Overview

This memo pre-supposes an operational mode of behavior along the lines of the following:

1) As a matter of network design, specific links (or engineered virtual links) are set aside for the purpose of acting as shared protection resources. The key attribute of these links is that the processing of TC markings will be exclusively for shared protection.

2) The arrangement of the shared protection links can be arbitrary such that contiguous domains can be constructed with an arbitrary number of ingress and egress points. A set of contiguous protection links is known as a protection domain.

3) Either an apriori or on-demand mesh of SPMEs that connect all ingress and egress points in a protection domain is required. These are logically forwarding adjacencies for the purposes of routing protection paths

4) The instantiation of protection paths requires the mapping of the incoming path at an ingress node for the protection domain to an SPME that connects the ingress to the required egress node from the domain.

5) The pipe model of TC copying is used such that the SPME gets the TC marking associated with the business priority for the path associated with the incoming label value. As the SPME only transits resources where the TC marking has been overloaded in this fashion business priority does not conflict with application requirements.

6) Admission control for the protection paths transiting the protection domain is performed such that the sum of the bandwidth for a given business priority does not oversubscribe any links in the protected domain, but the sum of the bandwidth for all business priorities can. In this way no traffic of the highest business priority using the shared mesh pool will be contended.

4. Signalling Implications

For a future version of this document.

5. IANA Considerations

No IETF protocols were harmed in the publishing of this memo.

6. Security Considerations

For a future version of this document.

7. References

7.1. Normative References

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

7.2. Informative References

[2] Sprecher, N., et al. "MPLS Transport Profile (MPLS-TP) Survivability Framework", RFC 6372, September 2011

8. Authors' Addresses

Dave Allan
Ericsson
Email: david.i.allan@ericsson.com

Greg Mirsky
Ericsson
Email: Gregory.mirsky@ericsson.com