Network Working Group
Internet-Draft
Intended status: Experimental
Expires: July 5, 2007

L. Andersson
A. Gavler
Acreo AB
January 2007

Extenstion to RSVP-TE for GMPLS Controlled Ethernet - An experimental approach
draft-andersson-gels-exp-rsvp-te-01.txt

Status of this Memo

Copyright Notice

Abstract

This document specifies the extensions to RSVP-TE that Acreo AB has used in the GMPLS part of testbed.

Table of Contents

1. Introduction

This Internet Draft documents the extensions to RSVP-TE that were used in the tests of GMPLS Controlled Ethernet (GELS), which were performed in the Acreo National Broadband Testbed end of 2006 and early 2007.

In Section 2 we give a short background of the research in the test bed in general and the GMPLS controlled Ethernet in particular.

Note: The -01 draft has been updated after comments on the gels mailing list, and does when it is written not exactly match our implementation.

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 [RFC2119].

2. GMPLS Controlled Ethernet research

2.1. The Acreo National Broadband Testbed

The Acreo National Broadband Testbed (ANBT) has been set up a joint effort by the Swedish government and the Swedish industry. Acreo AB was chosen to host the test bed, and the task is to initiate research projects on different aspects of broadband networks. Methods for control of carrier Ethernet has attracted quite a bit of interest.

GELS test bed is subset of the ABNT and consists of a 6 GMPLS enabled IP routers and 4 Ethernet Bridges. In some of our tests we've also used Linux based SW Ethernet Bridges.

2.2. Ethernet Control Plane

The control plane as implemented in the ANBT consists of three different parts:

- Routing Protocol - OSPF-TE, we are running the OSPF-TE exactly as implemented by the Dragon project.

- Signaling protocol - RSVP-TE, we have made the extensions to RSVP-TE, RFC3471 [RFC3471] and RFC3473 [RFC3473] as specified in Section 3.

- Link Management Protocol - LMP, we have not yet implemented the LMP protocol.

2.3. The Ethernet data plane

In the test bed we have used the Ethernet data plane as specified in IEEE Std 802.1Q. This has been made possible since we control the entire network by a GMPLS control plane and by default set up loop free LSPs. We have no traffic entering the network that results in that flooding or learning is triggered. This is clearly an artificial condition, but it is very well acceptable in a research network.

To take GELS into production networks is outside the scope of the current work we've undertaken, our focus is to establish a test bed e.g. for tests of new control plane extensions, traffic engineering paradigms and advanced applications. To run a GMPLS control plane for a production network will quite possibly require 802.1Q S-VLAN tags as specified in the IEEE Std 802.1ad Provider Bridging amendment to IEEE Std 802.1Q. and possibly the future IEEE802.1ah standard.

2.4. Motivation for a GMPLS controlled Ethernet

The answer to question "Why GELS?" is simple from a research perspective. Very much of research starts from the question "What happens if ...?"

In this case the question was "What happens if we make use of an GMPLS control plane to control an Ethernet network?" The answer to that question will decide whether we'll continue using GELS a as configuration tool while setting up tests in our network. Two tentative results today is that (1) for the application we have it is working well and saves us time, and (2) that we will look into the possibility to control every dynamic or configurable technology by the GMPLS control plane.

In addition to this we also have a number of external parties interested in GELS.

We have not been the only party active in this area, we have had a number of communications with e.g. Dave Allan, Don Fedyk, Dimitri Papadimitriou, Adrian Farrel, Attila Takacs, Deborah Brungard, Jai Hyung Cho and Nurit Sprecher. They have not reviewed this document, but nevertheless have had influence on our thinking on the subject. This is the major reason to share what we've been doing.

We are also in debt to the Dragon project, that gave us a good start when we could use their open source code as a starting point.

2.5. Incremental development

Our general approach to GELS has been stable over time, we've wanted to use the possibility to statically configure Ethernet Bridges by means of a GMPLS signalling protocol and to learn network topology and traffic engineering information by means of OSPF-TE.

One thing has been changing though; our understanding what the "Ethernet label" is and how it can be used.

- Our first approach was that it would be possible to define a new Tag Protocol Identifier (tpid) that should have pointed to "Ethernet Label" rather than a 802.1Q VLAN tag. Since this idea involved major changes to the Ethernet data plane, the Ethernet Standards and existing implementations this idea were quickly dropped. However we were able to prove that the concept works on off the shelf equipment.

- Our second approach was to simply use the 802.1Q VLAN tag, but due to scalability problems (4096 labels per network) we wanted to

swap them per link. The IEEE802.1 have made it very clear this also breaks existing IEEE standards. However, communications from IEEE802.1 have opened up for a certain type of VID swapping. There are indication that the idea of VID swapping, which is accepted at certain types of interface, might be increasingly accepted in the future.

- At this point a number of ideas started to emerge from a lot of different sources. Today we are convinced that an Ethernet tag should be possible to use as an Ethernet label, often in combination with the destination MAC Address. Further we are mostly looking to 802.1Q S-VLAN tags as defined in IEEE Std 802.1ad Provider Bridging amendment to IEEE Std 802.1Q. It is possible that when the IEEE Std 802.1ah is ready the new tag defined there will be possible to use.

- Our current network works with standard 802.1Q bridges.

3. Protocol extensions

Taking a starting point in RFC3471 [RFC3471] and RFC3473 [RFC3473] we have made the following extensions and adaptations to RSVP-TE.

3.1. Information in the Generalized Label Request Object

The required information to be carried by a PATH message in the Label Request Object is defined in RFC3471, and the format of the Label Request Object is defined in RFC3473 as follows:

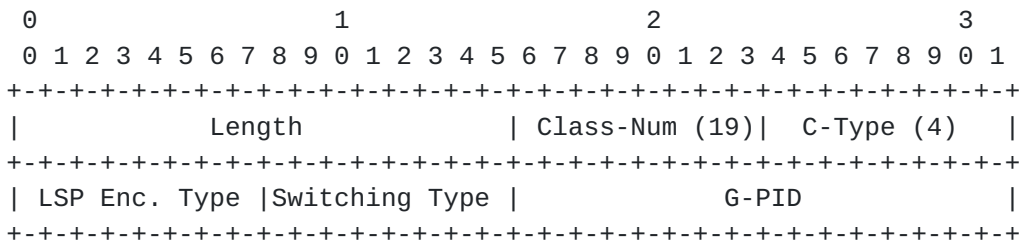

For the purpose of GELS we use the following encoding of the Label Request Object:

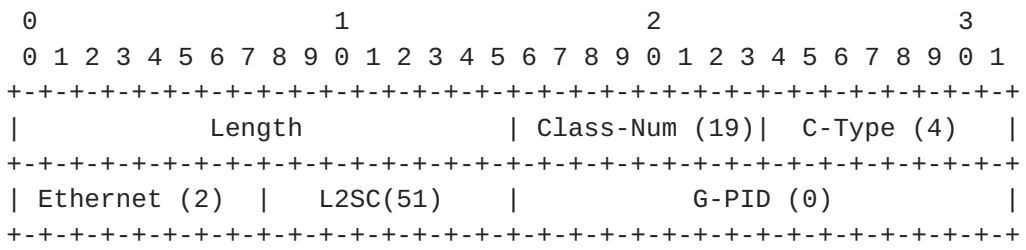

This is according to the parameter definitions in RFC3471.

3.2. Label Definition

The format of a Generalized Label object is:

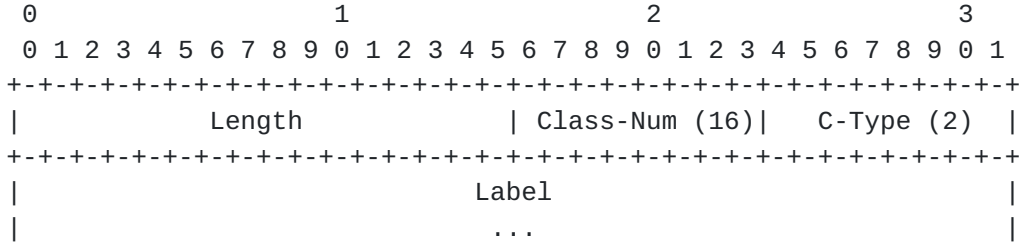

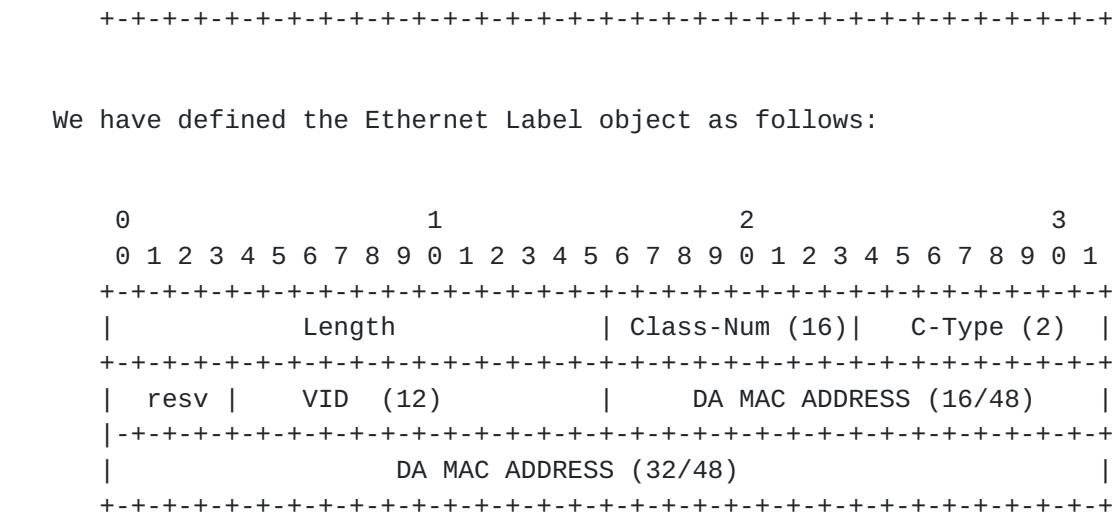

We have defined the Ethernet Label object as follows:

The four most significant bits of the Label field is not used and should be set to 0 went sent and ignored when received.

3.3. Extensions to the Session Attribute Object

The Session Attribute Object is optional and carried in the PATH message.

In RFC3209 [RFC3209] the Session Attribute Object is defined. The Session Attribute Class is 207. RFC3209 also defines two C_Types, LSP_TUNNEL, C-Type = 7 and LSP_TUNNEL_RA, C-Type = 1.

The LSP_TUNNEL_RA C-Type includes all the same fields as the LSP_TUNNEL C-Type. Additionally it carries resource affinity information. This document defines a third format LSP_TUNNEL_ETH, the C-type = 12. The LSP_TUNNEL_ETH C-type carries all the same fields as the LSP_TUNNEL_RA C-type. Additionally it carries Ethernet LSP attribute information.

We have defined the LSP_TUNNEL_ETH C-type as follows. The Session Attribute Class is 207 and the LSP_TUNNEL_ETH, the C-type is 12.

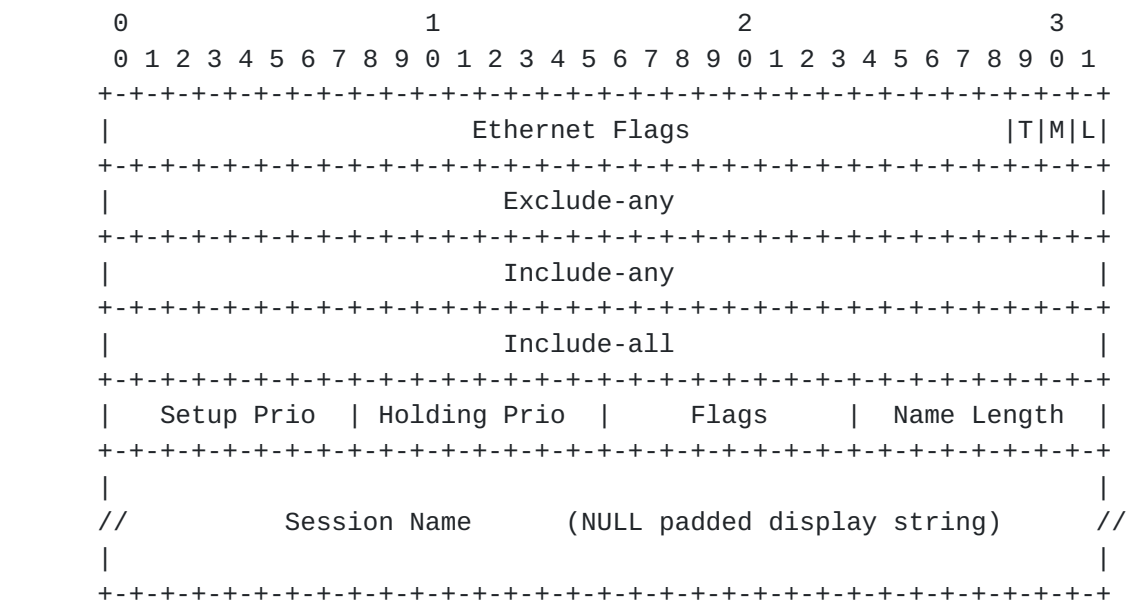

T = type bit, 1 indicates that both VID and DA MAC address is used, 0 indicates that only VID is used.

M = merge bit, 1 indicates that merging is allowed, 0 indicates that merging is not allowed.

L = learning bit, if the learning bit is set to 1 this indicates that GELS control plane is used to set up a standard IEEE802.1Q VLAN, i.e. learning, ageing, broadcast and a Multiple Spanning Tree Protocol (MSTP) will be turned on (L=1) or turned of (L=0).

bit 0 to 28 are reserved, and has to be set to 0 when sending and ignored when received.

3.4. Suggested Label Object

The suggested label object is optional and carried in the PATH message. The format of the suggest label is identical to the format of the Generalized Label object.

The information in the Suggested Label in combination with the flags in the LSP_TUNNEL_ETH C-type is interpreted as follows:

If the ingress node specifies a VID in the suggested label this is the VID to be used. If the VID field is set to all zeroes, this is an indication that no VID is specified.

The DA MAC Address field should always be set to all zeroes by the ingress LSR in a Suggested Label object, and the field SHALL always be ignored when received.

4. Procedures

When sending a PATH message the ingress LSR may include a Suggested Label object and/or a Session Atribute Object (C-num = 12). The information in the Suggested Label object and/or Session Attribute Object will be used by the nodes to determine the type of LSP requested.

If the ingress LSR does not include a Suggested Label object or a Serssion Attribute object in the PATH message, the egress LSR or merge LSR will treat it as if it were a request for an merge capable LSP with a label consisting of both a VID and a DA MAC address.

When an intermediate LSR receives a PATH message with a Suggested Label object and/or a Session Attribute objcet it MUST forward these objects unchanged, unless it is able to merge on to an existing LSP. The criteria for merging is for further study.

An egress LSR that receives a path message carrying a Label Request object but no Suggested Label object or any flags in the Session Attribute object WILL interpret this a a request for a merge capable LSP where both the VID and DA MAC Address is used as the label.

Note: This section may be extended.

5. Security Considerations

This document specify protocol extensions to RSVP-TE that is intended to be used in research contexts. Security consideration has therefore been left for further study and it is strongly recommended not to use these extensions in any network that is part of or connected to the Internet.

6. IANA Considerations

We will ask IANA to allocate C-type 12 for LSP_TUNNEL_ETH under the Session Attribute Class 207

7. Acknowledgements

-

8. References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC3473] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

Authors' Addresses

Loa Andersson
Acreo AB

Email: loa@pi.se

Anders Gavler
Acreo AB

Email: anders.gavler@acreo.se

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).