

Internet Engineering Task Force (IETF)
Internet-Draft
Intended status: Informational
Expires: April 7, 2016

A. Malis, Ed.
Huawei Technologies
B. Wilson
Applied Communication Sciences
G. Clapp
AT&T Labs Research
V. Shukla
Verizon Communications
October 5, 2015

Requirements for Very Fast Setup of GMPLS LSPs
draft-ietf-teas-fast-lsps-requirements-02

Abstract

Establishment and control of Label Switch Paths (LSPs) have become mainstream tools of commercial and government network providers. One of the elements of further evolving such networks is scaling their performance in terms of LSP bandwidth and traffic loads, LSP intensity (e.g., rate of LSP creation, deletion, and modification), LSP set up delay, quality of service differentiation, and different levels of resilience.

The goal of this document is to present target scaling objectives and the related protocol requirements for Generalized Multi-Protocol Label Switching (GMPLS).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on April 7, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Background	3
3.	Motivation	3
4.	Driving Applications and Their Requirements	5
4.1.	Key Application Requirements	5
5.	Requirements for Very Fast Setup of GMPLS LSPs	6
5.1.	Protocol and Procedure Requirements	6
6.	IANA Considerations	7
7.	Security Considerations	7
8.	Acknowledgements	7
9.	References	7
9.1.	Normative References	7
9.2.	Informative References	8
	Authors' Addresses	8

[1.](#) Introduction

Generalized Multi-Protocol Label Switching (GMPLS) [[RFC3471](#)] [[RFC3945](#)] includes an architecture and a set of control plane protocols that can be used to operate data networks ranging from packet-switch-capable networks, through those networks that use Time Division Multiplexing, to WDM networks. The Path Computation Element (PCE) architecture [[RFC4655](#)] defines functional components that can be used to compute and suggest appropriate paths in connection-oriented traffic-engineered networks. Additional wavelength switched optical networks (WSON) considerations were defined in [[RFC6163](#)].

This document refers to the same general framework and technologies, but adds requirements related to expediting LSP setup under heavy connection churn scenarios, while achieving low blocking under an overall distributed control plane. This document focuses on a

specific problem space - high capacity and highly dynamic connection request scenarios - that may require clarification and or extensions to current GMPLS protocols and procedures. In particular, the purpose of this document is to address the potential need for protocols and procedures that enable expediting the setup of LSPs in high churn scenarios. Both single-domain and multi-domain network scenarios are considered.

This document focuses on the following two topics: 1) the driving applications and main characteristics and requirements of this problem space, and 2) the key requirements which may be novel with respect to current GMPLS protocols.

This document presents the objectives and related requirements for GMPLS to provide the control for networks operating with such performance requirements. While specific deployment scenarios are considered as part of the presentation of objectives, the stated requirements are aimed at ensuring the control protocols are not the limiting factor in achieving a particular network's performance. Implementation dependencies are out of scope of this document.

It is envisioned that other documents may be needed to define how GMPLS protocols meet the requirements laid out in this document. Such future documents may define extensions, or simply clarify how existing mechanisms may be used to address the key requirements of highly dynamic networks.

2. Background

The Defense Advanced Research Projects Agency (DARPA) Core Optical Networks (CORONET) program [[Chiu](#)], is an example target environment that includes IP and optical commercial and government networks, with a focus on highly dynamic and resilient multi-terabit core networks. It anticipates the need for rapid (sub-second) setup and SONET/SDH-like restoration times for high-churn (up to tens of requests per second network-wide and holding times as short as one second) on-demand wavelength, sub-wavelength, and packet services for a variety of applications (e.g., grid computing, cloud computing, data visualization, fast data transfer, etc.). This must be done while meeting stringent call blocking requirements and while minimizing the use of resources such as time slots, switch ports, wavelength conversion, etc.

3. Motivation

The motivation for this document, and envisioned related future documents, is two-fold:

1. The anticipated need for rapid setup, while maintaining low blocking, of large bandwidth and highly churned on-demand connections (in the form of sub-wavelengths, e.g., OTN ODUx, and wavelengths, e.g., OTN OCh) for a variety of applications including grid computing, cloud computing, data visualization, and intra- and inter-datacenter communications.
2. The ability to set up circuit-like LSPs for large bandwidth flows with low setup delays provides an alternative to packet-based solutions implemented over static circuits that may require tying up more expensive and power-consuming resources (e.g., router ports). Reducing the LSP setup delay will reduce the minimum bandwidth threshold at which a GMPLS circuit approach is preferred over a layer 3 (e.g., IP) approach. Dynamic circuit and virtual circuit switching intrinsically provide guaranteed bandwidth, guaranteed low-latency and jitter, and faster restoration, all of which are very hard to provide in packet-only networks. Again, a key element in achieving these benefits is enabling the fastest possible circuit setup times.

Future applications are expected to require setup times as fast as 100 ms in highly dynamic, national-scale network environments while meeting stringent blocking requirements and minimizing the use of resources such as switch ports, wavelength converters/regenerators, and other network design parameters. Of course, the benefits of low setup delay diminish for connections with long holding times. The need for rapid setup for specific applications may override and thus get traded off, for these specific applications, against some other features currently provided in GMPLS, e.g., robustness against setup errors.

With the advent of data centers, cloud computing, video, gaming, mobile and other broadband applications, it is anticipated that connection request rates may increase, even for connections with longer holding times, either during limited time periods (such as during the restoration from a data center failure) or over the longer term, to the point where the current GMPLS procedures of path computation/selection and resource allocation may not be timely, thus leading to increased blocking or increased resource cost. Thus, extensions of GMPLS signaling and routing protocols (e.g., OSPF-TE) may also be needed to address heavy churn of connection requests (i.e., high connection request arrival rate) in networks with high traffic loads, even for connections with relatively longer holding times.

4. Driving Applications and Their Requirements

There are several emerging applications that fall under the problem space addressed here in several service areas such as provided by telecommunication carriers, government networks, enterprise networks, content providers, and cloud providers. Such applications include research and education networks/grid computing, and cloud computing. Detailing and standardizing protocols to address these applications will expedite the transition to commercial deployment.

In the target environment there are multiple Bandwidth-on-Demand service requests per second, such as might arise as cloud services proliferate. It includes dynamic services with connection setup requirements that range from seconds to milliseconds. The aggregate traffic demand, which is composed of both packet (IP) and circuit (wavelength and sub-wavelength) services, represents a five to twenty-fold increase over today's traffic levels for the largest of any individual carrier. Thus, the aggressive requirements must be met with solutions that are scalable, cost effective, and power efficient, while providing the desired quality of service (QoS).

4.1. Key Application Requirements

There are two key performance-scaling requirements in the target environment that are the main drivers behind this draft:

1. Connection request rates ranging from a few requests per second for high capacity (e.g., 40 Gb/s , 100 Gb/s) wavelength-based LSPs to around 100 requests per second for sub-wavelength LSPs (e.g., OTN ODU0, ODU1, and ODU2).
2. Connection setup delay of around 100 ms across a national or regional network. To meet this target, and assuming pipelined cross-connection, and worst case propagation delay and hop count, it is estimated that the maximum processing delay per hop is around 700 microseconds [[Lehmen](#)]. Optimal path selection and resource allocation may require somewhat longer processing (up to 5 milliseconds) in either the destination or source nodes and possibly tighter processing delays (around 500 microseconds) in intermediate nodes.

The model for a national network is that of the continental US with up to 100 nodes and LSPs distances up to ~3000 km and up to 15 hops.

A connection setup delay is defined here as the time between the arrival of a connection request at an ingress edge switch - or more generally a Label Switch Router (LSR) - and the time at which information can start flowing from that ingress switch over that

connection. Note that this definition is more inclusive than the LSP setup time defined in [[RFC5814](#)] and [[RFC6777](#)], which do not include PCE path computation delays.

5. Requirements for Very Fast Setup of GMPLS LSPs

This section lists the protocol requirements for very fast setup of GMPLS LSPs in order to adequately support the service characteristics described in the previous sections. These requirements may be the basis for future documents, some of which may be simply informational, while others may describe specific GMPLS protocol extensions. While some of these requirements may have implications on implementations, the intent is for the requirements to apply to GMPLS protocols and their standardized mechanisms.

5.1. Protocol and Procedure Requirements

- R1 The protocol processing related portion of the LSP establishment time should scale linearly based on number of traversed nodes.
- R2 End-to-end LSP data path availability should be bounded by the worst case single node data path establishment time. In other words, pipelined cross-connect processing as discussed in [[RFC6383](#)] should be enabled.
- R3 LSP Establishment time shall depend on the number of nodes supporting an LSP and link propagation delays and not any off (control) path transactions, e.g., PCC-PCE and PCC-PCC communications at the time of connection setup, even when PCE-based approaches are used.
- R4 Must support LSP holding times as short as one second.
- R5 The protocol aspects of LSP signaling must not preclude LSP request rates of tens per second.
- R6 The above requirements should be met even when there are failures in connection establishment, i.e., LSPs should be established faster than when crank-back is used.
- R7 These requirements are applicable even when an LSP crosses one or more administrative domains/boundaries.
- R8 The above are additional requirements and do not replace existing requirements, e.g., alarm free setup and teardown, Recovery, or inter-domain confidentiality.

6. IANA Considerations

This memo includes no requests to IANA.

7. Security Considerations

Being able to support very fast setup and a high churn rate of GMPLS LSPs is not expected to adversely affect the underlying security issues associated with existing GMPLS signaling. If encryption that requires key exchange is intended to be used on the signaled LSPs, then this requirement needs to be included as a part of the protocol design process, as the usual extra round trip time (RTT) for key exchange will have an effect on the setup and churn rate of the GMPLS LSPs. It is possible to amortize the costs of key exchange over multiple exchanges (if those occur between the same peers) so that some exchanges need not cost a full RTT and operate in so-called zero-RTT mode.

8. Acknowledgements

The authors would like to thank Ann Von Lehmen, Joe Gannett, Ron Skoog, and Haim Kobrinski of Applied Communication Sciences for their comments and assistance on this document. Lou Berger provided editorial comments on this document.

9. References

9.1. Normative References

- [RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", [RFC 3471](#), DOI 10.17487/RFC3471, January 2003, <<http://www.rfc-editor.org/info/rfc3471>>.
- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", [RFC 3945](#), DOI 10.17487/RFC3945, October 2004, <<http://www.rfc-editor.org/info/rfc3945>>.
- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), DOI 10.17487/RFC4655, August 2006, <<http://www.rfc-editor.org/info/rfc4655>>.
- [RFC5814] Sun, W., Ed. and G. Zhang, Ed., "Label Switched Path (LSP) Dynamic Provisioning Performance Metrics in Generalized MPLS Networks", [RFC 5814](#), DOI 10.17487/RFC5814, March 2010, <<http://www.rfc-editor.org/info/rfc5814>>.

- [RFC6163] Lee, Y., Ed., Bernstein, G., Ed., and W. Imajuku, "Framework for GMPLS and Path Computation Element (PCE) Control of Wavelength Switched Optical Networks (WSONs)", [RFC 6163](#), DOI 10.17487/RFC6163, April 2011, <<http://www.rfc-editor.org/info/rfc6163>>.
- [RFC6383] Shiomoto, K. and A. Farrel, "Advice on When It Is Safe to Start Sending Data on Label Switched Paths Established Using RSVP-TE", [RFC 6383](#), DOI 10.17487/RFC6383, September 2011, <<http://www.rfc-editor.org/info/rfc6383>>.
- [RFC6777] Sun, W., Ed., Zhang, G., Ed., Gao, J., Xie, G., and R. Papneja, "Label Switched Path (LSP) Data Path Delay Metrics in Generalized MPLS and MPLS Traffic Engineering (MPLS-TE) Networks", [RFC 6777](#), DOI 10.17487/RFC6777, November 2012, <<http://www.rfc-editor.org/info/rfc6777>>.

[9.2.](#) Informative References

- [Chiu] A. Chiu, et al, "Architectures and Protocols for Capacity Efficient, Highly Dynamic and Highly Resilient Core Networks", Journal of Optical Communications and Networking vol. 4, No. 1, pp. 1-14, January 2012, <<http://dx.doi.org/10.1364/JOCN.4.000001>>.
- [Lehmen] A. Von Lehmen, et al, "CORONET: Testbeds, Demonstration and Lessons Learned", Journal of Optical Communications and Networking Vol. 7, Issue 3, pp. A447-A458, March 2015, <<http://www.osapublishing.org/jocn/abstract.cfm?uri=jocn-7-3-A447>>.

Authors' Addresses

Andrew G. Malis (editor)
Huawei Technologies

Email: agmalis@gmail.com

Brian J. Wilson
Applied Communication Sciences

Email: bwilson@appcomsci.com

George Clapp
AT&T Labs Research

Email: clapp@research.att.com

Vishnu Shukla
Verizon Communications

Email: vishnu.shukla@verizon.com