

Network Working Group
Internet Draft
Intended Status: Informational
Expires: April 14, 2013

X. Fu(Ed.), M. Betts, Q.
Wang
ZTE
V. Manral
Hewlett-Packard Corp.
D. McDysan(Ed.), A. Malis
Verizon
S. Giacalone
Thomson Reuters
J. Drake
Juniper Networks

October 15, 2012

Delay and Loss Traffic Engineering Problem Statement for MPLS

[draft-fuxh-mpls-delay-loss-te-problem-statement-01](#)

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 February 27, 2013.

Copyright Notice

Copyright (c) 2012 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.

McDysan

Expires April 14, 2013

[Page 1]

Abstract

Deployment and usage of cloud based applications and services that use an underlying MPLS network are expanding and an increasing number of applications are extremely sensitive to delay and packet loss. Furthermore, in cloud computing an additional decision problem arises of simultaneously choosing the data center to host applications along with MPLS network connectivity such that the overall performance of the application is met. Mechanisms exist to measure and monitor MPLS path performance parameters for packet loss and delay, but the mechanisms work only after the path has been setup. The cloud-based and performance sensitive applications would benefit from measurement of MPLS network and potential path information that would be provided for use in the computation before LSP setup and then the selection of LSPs.

This document provides a statement of problems faced by these cloud based and performance sensitive applications and describes requirements to enable the efficient and accurate measurement of the MPLS network. This also allows new performance parameters to be reported and used in the computation of MPLS services in support of these cloud based and performance sensitive applications.

Table of Contents

<u>1.</u>	<u>Introduction.....</u>	<u>3</u>
<u>1.1.</u>	<u>Scope.....</u>	<u>3</u>
<u>2.</u>	<u>Conventions used in this document.....</u>	<u>3</u>
<u>2.1.</u>	<u>Acronyms.....</u>	<u>3</u>
<u>2.2.</u>	<u>Terminology and Assumptions.....</u>	<u>4</u>
<u>2.2.1.</u>	<u>Delay.....</u>	<u>4</u>
<u>2.2.2.</u>	<u>Packet Loss.....</u>	<u>4</u>
<u>2.2.3.</u>	<u>Packet Delay Variation.....</u>	<u>5</u>
<u>3.</u>	<u>Motivation and Background.....</u>	<u>5</u>
<u>3.1.</u>	<u>General Characteristics of Performance Parameters.....</u>	<u>5</u>
<u>3.2.</u>	<u>Use Cases for Performance Parameter Sensitive LSP Placement</u>	<u>5</u>
<u>4.</u>	<u>Problem Statement.....</u>	<u>6</u>
<u>4.1.</u>	<u>End-to-end Measurement Insufficient for Performance Sensitive LSP Path Selection.....</u>	<u>6</u>
<u>4.2.</u>	<u>Lower Layer MPLS Networks Unable to Communicate Significant Performance Changes.....</u>	<u>7</u>
<u>4.3.</u>	<u>No Method to Communicate Significant Node/Link Performance Changes.....</u>	<u>7</u>
<u>4.4.</u>	<u>Routing Metrics Insufficient for Performance Sensitive Path Selection.....</u>	<u>7</u>
<u>4.5.</u>	<u>LSP Signaling Methods Insufficient for Performance Sensitive Path Selection.....</u>	<u>8</u>
<u>5.</u>	<u>Functional Requirements.....</u>	<u>8</u>
<u>5.1.</u>	<u>Augment LSP Requestor Signaling with Performance Parameter Values.....</u>	<u>8</u>
<u>5.2.</u>	<u>Specify Criteria for Node and Link Performance Parameter Estimation, Measurement Methods.....</u>	<u>9</u>

5.3. Support Node Level Performance Information when Needed...9
5.4. Augment Routing Information with Performance Parameter Estimates
.....10
5.5. Augment Signaling Information with Concatenated Estimates10

5.6. Define Significant Performance Parameter Change Thresholds and Frequency.....	<u>10</u>
<u>5.7.</u> Define Thresholds and Timers for Links with Unusable Performance	<u>11</u>
5.8. Communicate Significant Performance Changes between Layers	11
5.9. The above requirement applies to layering with different technologies (e.g., MPLS over OTN) or to different levels within the same technology (e.g., hierarchical LSPs).....	<u>11</u>
<u>5.10.</u> Support for Networks with Composite Link.....	<u>11</u>
<u>5.11.</u> Restoration, Protection and Rerouting.....	<u>11</u>
<u>5.12.</u> Management and Operational Requirements.....	<u>12</u>
<u>6.</u> IANA Considerations.....	<u>12</u>
<u>7.</u> Security Considerations.....	<u>12</u>
<u>8.</u> References.....	<u>12</u>
<u>8.1.</u> Normative References.....	<u>12</u>
<u>8.2.</u> Informative References.....	<u>12</u>
<u>9.</u> Acknowledgments.....	<u>13</u>

1. Introduction

This draft is one of two created from [draft-fuxh-mpls-delay-loss-te-framework-05](#) in response to comments from an MPLS Review Team (RT). This draft focuses on a problem statement and requirements, for delay and loss based Traffic Engineering for MPLS networks. A peer document focuses on the framework.

The intent of this document is to focus on stating the technical aspects of the application oriented problems to be solved and specific requirements targeted to solve these problems.

It describes requirements and application needs for bounded values of delay, packet loss and delay variation.

1.1. Scope

A (G)MPLS network may have multiple layers of packet, TDM and/or optical network technology and an important objective is to make a prediction of end-to-end delay, loss and delay variation based upon the current state of this network with acceptable accuracy before an LSP is established.

The (G)MPLS network may cover a single IGP area/level, may be a hierarchical IGP under control of a single administrator, or may involve multiple domains under control of multiple administrators.

2. Conventions used in this document

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](#)].

2.1. Acronyms

SLA Service Level Agreement

SLS Service Level Specification

Fu et al

Expires April 14, 2013

[Page 3]

NPO Network Performance Objective

2.2. Terminology and Assumptions

A Service Level Agreement (SLA) is a contractual agreement that service providers have with customers for services comprised of numerical values for performance measures; for example, delay, loss and delay variation. Additionally, network operators may have Service Level Specification (SLS) that is for internal use by the operator. See [ITU-T.Y.1540], [ITU-T.Y.1541], RFC 3809, Section 4.9 [RFC3809] for examples of the form of such SLA and SLS specifications.

Network Performance Objective (NPO) is defined in section 5 of [ITU-T.Y.1541] in terms of numerical values for performance measures, principally delay, loss, and delay variation. The term NPO is used in this document since the SLA and SLS measures have network operator and service specific implications. Furthermore, the NPO measures are sufficiently well defined to address other use cases and the stated problems.

Of particular interest is the composition methods defined in Y.1541 for estimating performance parameters of candidate LSP paths based upon the performance parameter estimates/measurements of individual nodes and links.

This document assumes that the evaluation interval for a performance parameter is on the order of minutes as stated in [ITU-T Y.1541, 5.3.2], which is the same of that used in some commercial networks.

2.2.1. Delay

Section 6.2.1 of [ITU-T Y.1540] defines mean IP Packet Transfer Delay (IPTD) as the arithmetic average of the one-way delay observed between measurement points IPTD is referred to as "delay" in this document.

Section 8.2.1 of [ITU-T Y.1541] defines composition of the IPTD UNI-UNI performance parameter as "the mean IP packet transfer delay (IPTD) performance parameter, the UNI-UNI performance is the sum of the means contributed by network sections."

2.2.2. Packet Loss

Section 6.4 of [ITU-T Y.1540] defines IP Packet Loss Ratio (IPLR) as the "ratio of total lost IP packet outcomes to total transmitted IP packets in a population of interest," which is referred to as "loss" in this document.

Section 8.2.2 of [ITU-T Y.1541] defines composition of the IPLR UNI-UNI performance parameter as "may be estimated by inverting the probability of successful packet transfer across n network sections."

2.2.3. Packet Delay Variation

Section 6.2.4.2 of [ITU-T Y.1540] defines quantile-based limits on IP Packet Delay Variation (IPDV), which is referred to as "delay variation" in this document.

Section 8.2.4 of [ITU-T Y.1541] defines composition of the IPDV UNI-UNI performance parameter as "must recognize their sub-additive nature and it is difficult to estimate accurately without considerable information about the individual delay distributions." Appendix IV of [ITU-T Y.1541] gives several examples of IPDV estimate calculations.

3. Motivation and Background

3.1. General Characteristics of Performance Parameters

In general, nodes and links contribute to the performance parameters in the network. Another significant contributor is that of the host stack, but that is outside the scope of this document.

For many applications, the delay NPO is very important. In networks with wide geographic separation, propagation delay may dominate delay, while in local or metro networks nodal delay may become important.

Some link technologies (e.g., wireless, wifi, satellite) may have packet loss characteristics inherently different from those of other link technologies (e.g., fiber optic, cable) networks. Packet loss can be caused due to signal degradation/ high noise, which causes corrupted packets which in turn are discarded in the network. Furthermore, the loading of queues (congestion) may also result in packet loss.

Delay variation (sometimes also referred to as packet jitter) is important to some applications, such as interactive voice, video and/or multimedia communication, gaming, and simulations. If delay varies too much, then a playback buffer for such applications may underflow or overflow, resulting in a disruption to the application. Delay variation is caused primarily by queuing within a node. It can also be caused when packets take different paths, due to lower layer routing.

3.2. Use Cases for Performance Parameter Sensitive LSP Placement

In High Frequency trading for Electronic Financial markets, computers make decisions based on the Electronic Data received, without human intervention. These trades now account for a majority of the trading volumes and rely exclusively on ultra-low-delay direct market access. In certain networks, such as financial information networks (e.g. stock market data providers), network performance information (e.g. delay) is critical to data path selection. In these networks, extremely large amounts of money rest on the ability to access market data as quickly as possible and to predictably make trades faster than the competition. Using metrics such as hop count or link cost as routing may not always meet this need. In such networks it would be beneficial to be able to

make path selection decisions based on performance data (such as delay)
in a cost-effective and scalable way.

In other networks, for example, network-based VPNs there are in place between a customer and a provider a Service Level Agreement (SLA) which specifies performance objectives, such as delay, loss, and delay variation. In some cases these performance objectives are defined between specific customer locations. Furthermore, packets may be associated with certain classes as identified by packet header fields (e.g., IP DSCP, IEEE P-bits, MPLS TC bits) that are associated with different performance objectives. In these types of networks, the objective is to provide service that is no worse than the performance objective. A single SLA may support many customers of the same type. There is also a need to support specific SLAs, typically for very large customers who demand premium performance for which they are willing to pay a premium price.

In emerging cloud-based services, an additional decision problem where the application may be placed in a choice of more than one data center and the (G)MPLS network connectivity may also be chosen [[CLO](#), [CSO](#)]. In these types of applications, the objective is to meet the overall performance of the application deployed in one more or more data centers. The performance of the intra- data center performance component is out of scope of this draft, but this overall cloud plus networking decision problem would benefit from a prediction of the MPLS network performance as part of path establishment.

4. Problem Statement

With the use cases in the previous section as motivation, there are several technical problems that currently standardized IETF protocols do not adequately address:

- o End-to-end Measurement Insufficient for Performance Sensitive LSP Path
- o Routing Metrics Insufficient for Performance Sensitive Path Selection
- o LSP Signaling Methods Insufficient for Performance Sensitive Path Selection
- o Lower Layer MPLS Networks Unable to Communicate Significant Performance Changes
- o No Method to Communicate Significant Node/Link Performance Changes

The following sections expand on each of these technical problem areas in more detail. Although some of the problem statements are made in terms of existing/proposed protocols, there is no intention to imply that the solution requires a revision to these protocols.

4.1. End-to-end Measurement Insufficient for Performance Sensitive LSP Path Selection

Methods exist to measure established LSP performance, e.g., [[RFC 6374](#)]

for MPLS-TP, and are most useful in verifying support for an NPO. [RFC 6374](#) specifies a mechanism to measure and monitor performance parameters for packet loss, and one-way and two-way delay, delay variation and

throughput. However, if measured performance is not met for an LSP there is not a standardized method to aid in an LSP originator or a proxy (e.g., PCE) to select a modified path that would meet the performance objective.

Therefore, there is a need to enable path computation that has access to an up to date recent performance estimate.

4.2. Lower Layer MPLS Networks Unable to Communicate Significant Performance Changes

Historically, when an IP/MPLS network was operated over a lower layer circuit switched network (e.g., SONET rings), a change in delay caused by the lower layer network (e.g., due to a maintenance action or failure) this was not known to the MPLS network. This resulted in delay affecting end user experience, sometimes violating NPO, SLS and/or SLA values and/or resulting in user complaints.

Using lower layer networks to provide restoration and grooming may be more efficient than performing packet only restoration, but the inability to communicate performance parameters, in particular delay, from the lower layer network to the higher layer network is an important problem to be solved in not only the composite link case [CL-REQ, [section 4.2](#)], but also in the case of single links connecting nodes.

In summary, Multi-layer GMPLS networks do not have a means to communicate a significant change in performance (e.g., delay) from one layer to another.

4.3. No Method to Communicate Significant Node/Link Performance Changes

Performance characteristics of links and nodes may change dynamically in response to a number of events. There is currently no way to automatically indicate which nodes and/or links have had significant performance changes to LSP originators or proxies so that they can attempt to recompute and signal a path that would meet the LSP performance objective.

4.4. Routing Metrics Insufficient for Performance Sensitive Path Selection

Optimization on a single metric does not meet the needs for all cases of performance sensitive path selection. In some cases, minimizing delay relates directly to the best customer experience (e.g., in TCP closer is faster or in financial trading the absolute minimum delay possible provides a competitive advantage). In other cases, user experience is relatively insensitive to delay, up to a specific limit at which point user perception of quality degrades significantly (e.g., interactive human voice and multimedia conferencing). A number of NPOs have a bound on point-point delay, and as long as this bound is met, the NPO is met - - decreasing the delay is not necessary. In some NPOs, if the specified delay is not met, the user considers the service as unavailable. An unprotected LSP can be manually provisioned on a set of links to meet

this type of NPO, but this lowers availability since an alternate route that meets the delay NPO cannot be determined.

One operational approach is to provision IP/MPLS networks over unprotected circuits and set the metric and/or TE-metric proportional to delay. This resulted in traffic being directed over the least delay path, even if this was not needed to meet an NPO or user experience objectives. This results in reduced flexibility and increased cost for network operators. However, the (TE) metric is often used to represent other information, such as link speed, economic cost or in support of ECMP (as described below) and may not be able to be set to be proportional to delay. Furthermore, if performance metrics such as loss and delay variation are to be supported in path selection, then proportional mapping is not possible.

Link attributes and LSP affinities [RFC 3209] can be used operationally to encode some information regarding performance, for example, indicating wired versus wireless, satellite versus terrestrial, etc. However, these attributes/affinities are used to encode other attributes and the 32 bit format is limiting in terms of numerical representation of performance objective parameters.

Another operational approach is to set (TE) metrics to (nearly) the same value so that LSPs are placed across multiple links using Equal Cost Multi-Path (ECMP) path selection. However, these parallel links may have markedly different performance characteristics (e.g., delay) and choice of a link that meets the performance objective is needed [CL-REQ, section 4.3].

IGP link and TE metrics are not sufficient to support performance sensitive path selection in a single IGP area/level [EXPRESS-PATH].

4.5. LSP Signaling Methods Insufficient for Performance Sensitive Path Selection

Current signaling approaches do not support inter area/level or inter-domain performance sensitive path selection. There is no standard for setting link attributes and LSP resource affinities [RFC 3209] between administrative domains, and since these have been used within some domains they are not a viable candidate to solve the aforementioned problems in this context. Augmenting an IGP with performance information does not solve the problem in these cases.

What is needed is a means for the originator/proxy of an LSP to confirm whether the estimated performance of a computed LSP path will meet the performance objective.

5. Functional Requirements

This section groups functional requirements intended to address the problems stated in the previous section into related areas.

5.1. Augment LSP Requestor Signaling with Performance Parameter Values

The solution needs to provide a means for an LSP requestor to signal

performance parameter sensitive paths. The following requirements state the types of requests that are required.

The solution MUST provide a means to indicate which performance parameters are supported by the network area/level or domain.

The solution MUST provide a means for the LSP requestor to ask for the minimum possible value for each supported performance parameter.

For example, an LSP requestor may ask for an LSP that has the minimum possible value of delay.

The solution MUST provide a means for the LSP requestor to ask for a range of acceptable values for each supported performance parameter.

For example, an LSP requestor may ask for an LSP that has performance between a minimum value of delay and packet loss and a maximum value of delay and packet.

5.2. Specify Criteria for Node and Link Performance Parameter Estimation, Measurement Methods

The solution MUST provide a means to configure the one-way link and node performance parameters for delay, loss and delay variation.

The solution SHOULD provide a means to dynamically measure and/or estimate the one-way link and node performance parameters for delay, loss and delay variation.

As defined in [section 2.2.](#) , the estimation interval for the performance parameters is assumed to be on the order of minutes. The solution MUST not impact stability nor significantly increase convergence time if performance parameters change over a timeframe on the order of minutes.

5.3. Support Node Level Performance Information when Needed

There are several scenarios under which node-related performance parameters (delay, loss, delay variation) has a different level of importance:

1. The case of few nodes with large geographic separation, (e.g., trans-oceanic), where link delay alone would be a good approximation.
2. The case of many nodes with small geographic separation (e.g., interconnected nearby data centers) where node/device delay is very important but link delay may be negligible.
3. The case of some number of nodes with medium geographic separation, where usage of both link and node delay may be desirable.

The intent in case 1 is to measure the predominant delay in uncongested service provider networks, where geographic delay dominates and is on the order of milliseconds or more. The argument in cases 2 and 3 for including node-level queuing performance parameters is that it better represents the performance experienced by applications. The argument

against including queuing related performance parameters is that if used in routing decisions it can result in routing instability. This tradeoff is discussed in detail in [CL-FW, [Section 4.1.1](#)].

The solution MUST define methods to include node level performance estimate information to routing protocols.

The solution MUST define methods to include node level performance estimate information to signaling protocols.

A specific deployment of the solution MAY choose to not use the node level performance estimates.

5.4. Augment Routing Information with Performance Parameter Estimates

The solution MUST provide a means to communicate performance parameters of both links and nodes as an estimate for use in performance sensitive LSP path selection within nodes of a single IGP area/level.

The solution SHOULD provide a means to communicate delay, loss and delay variation of links and nodes as a traffic engineering performance parameter for use in performance sensitive LSP path selection across a set of nodes in a hierarchy of IGP areas/levels.

5.5. Augment Signaling Information with Concatenated Estimates

The solution MUST provide a means to signal concatenated performance parameter estimates for both links and nodes as an estimate for use in performance sensitive LSP path selection traversing two or more separate administrative domains. See the terminology section for references on the concatenation method for specific performance parameters.

For example, the solution needs to support the capability to compute a route with X amount of bandwidth with less than Y ms of delay and less than Z% loss across multiple domains.

The solution MUST support the means to concatenate performance parameter estimates and report this for each traversed domain on the end-end path

The solution MUST interoperate with existing path selection and signaling methods traversing multiple domains.

5.6. Define Significant Performance Parameter Change Thresholds and Frequency

Delay, loss and delay variation measurements and/or estimates may be time varying. The solution MUST provide a means to control the advertisement rate of performance parameter estimates to avoid instability.

Any automatic LSP routing and/or load balancing solutions MUST NOT oscillate such that performance observed by users changes such that an NPO is violated. Since oscillation may cause reordering, there MUST be means to control the frequency of changing the path over which an LSP is placed.

5.7. Define Thresholds and Timers for Links with Unusable Performance

The solution MUST provide a means to configure a performance parameter threshold which defines placement of a node or link into an unusable state. The solution MUST provide a means to configure a performance parameter threshold which defines transition of a node or a link from an unusable state to a useable state. The solution MUST provide a means to control the minimum transition time between these states.

This unusable state is intended to operate on a link/node capability basis and not a global basis. Since state transition conditions are locally configured, all routers within a domain should synchronize this configuration value.

With current TE protocols, a refreshed LSP would use the most recent performance parameter estimates and may be rerouted based upon nodes or links being placed in an unusable performance state. [Section 5.11](#). defines requirements for a desirable function where performance sensitive LSP re-routing would occur.

5.8. Communicate Significant Performance Changes between Layers

In order to support network NPOs and provide acceptable user experience, the solution MUST specify a protocol means to allow a lower layer server network to communicate performance parameters (e.g., delay, loss, delay variation) to the higher layer client network.

5.9. The above requirement applies to layering with different technologies (e.g., MPLS over OTN) or to different levels within the same technology (e.g., hierarchical LSPs).

5.10. Support for Networks with Composite Link

An LSP may traverse a network with Composite Links [CL-REQ]. The solution's selection of performance sensitive paths SHOULD be compatible with the general availability, stability and transient response requirements of [CL-REQ, [Section 4.1](#)].

When an LSP traverses a network with composite links that has component links provided by lower layer networks, the solution MUST interoperate with the requirements [CL-REQ, [Section 4.2](#)].

When an LSP traverses a network with composite links that has parallel component links with different characteristics, the solution MUST interoperate with the requirements [CL-REQ, [Section 4.3](#)].

5.11. Restoration, Protection and Rerouting

The ability to re-route an LSP if one or more NPO objectives are not met is highly desirable. The solution SHOULD support the capability to configure an LSP as capable of implementing performance sensitive re-routing, as detailed in the following conditional requirements.

If performance sensitive re-routing is implemented, the solution MUST provide a means to configure performance parameter threshold crossing and time values.

If performance sensitive re-routing is implemented, the solution MUST support a configuration option to move an end-to-end LSP away from any link or node whose performance violates the configured threshold.

If implemented, the solution MUST provide a means to control the frequency of LSP rerouting to avoid instability.

If performance sensitive re-routing is implemented, and revertive behavior to a preferred LSP is supported, then the preferred LSP MUST not be released. When the end-to-end performance of the preferred LSP becomes acceptable, the service is restored to this preferred LSP.

The delay performance of pre-defined protection or dynamic reroutable LSP MUST be defined by the solution in terms of the maximum acceptable delay difference between the primary and protection/restoration path MUST be specifiable in the solution. For example, [MPLS-TP-USE-CASE] defines a Relative Delay Time which is the difference of the Absolute Delay between the primary and protection path.

5.12. Management and Operational Requirements

Existing management and diagnostic protocols MUST be able to operate over networks supporting performance sensitive LSP placement.

If performance sensitive re-routing is implemented, and end-to-end measurements of the LSP performance are made, then the LSP requestor is able to request path placement for a performance sensitive LSP using the previously stated requirements. Since a threshold crossing of the end-to-end performance measurement may or may not correspond to a change in the concatenated performance parameter estimates, making any automatic decision on this basis MUST not create instability.

6. IANA Considerations

No new IANA consideration are raised by this document.

7. Security Considerations

This document raises no new security issues.

8. References

8.1. Normative References

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

8.2. Informative References

[CL-UC] C. Villamizer et al, "Composite Link Use Cases and Design Considerations," [draft-ietf-rtgwg-cl-use-cases-01](#)

Fu et al

Expires April 14, 2013

[Page 12]

- [CL-REQ] C. Villamizar et al, "Requirements for MPLS Over a Composite Link", [draft-ietf-rtgwg-cl-requirement-08](#) .
- [CL-FW] C. Villamizar et al, "Composite Link Framework in Multi Protocol Label Switching (MPLS)", work in progress
- [ITU-T.Y.1540] ITU-T, "Internet protocol data communication service - IP packet transfer and availability performance parameters", 2011, <<http://www.itu.int/rec/T-REC-Y.1540/en>>.
- [ITU-T.Y.1541] ITU-T, "Network performance objectives for IP-based services", 2011, <<http://www.itu.int/rec/T-REC-Y.1541/en>>.
- [RFC3809] Nagarajan, A., "Generic Requirements for Provider Provisioned Virtual Private Networks (PPVPN)", [RFC 3809](#), June 2004.
- [CLO] Young Lee et al, "Problem Statement for Cross-Layer Optimization," work in progress.
- [CS0] Greg Bernstein, Young Lee, "Cross Stratum Optimization Use-cases," work in progress.
- [EXPRESS-PATH] A. Atlas, "Performance-based Path Selection for Explicitly Routed LSPs", work in progress.
- [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.
- [MPLS-TP-USE-CASE] L. Fang, "MPLS-TP Applicability; Use Cases and Design", [draft-ietf-mpls-tp-use-cases-and-design-01](#) .

9. Acknowledgments

This document was prepared using 2-Word-v2.0.template.dot.

The authors would like to thank the MPLS Review Team of Stewart Bryant, Daniel King and He Jia for their many helpful comments suggestions in July 2012.

Copyright (c) 2012 IETF Trust and the persons identified as authors of the code. All rights reserved.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING

NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This code was derived from IETF RFC [insert RFC number]. Please reproduce this note if possible.

Authors' Addresses

Xihua Fu
ZTE
Email: fu.xihua@zte.com.cn

Vishwas Manral
Hewlett-Packard Corp.
191111 Pruneridge Ave.
Cupertino, CA 95014
US
Phone: 408-447-1497
Email: vishwas.manral@hp.com

Dave McDysan
Verizon
Email: dave.mcdysan@verizon.com

Andrew Malis
Verizon
Email: andrew.g.malis@verizon.com

Spencer Giacalone
Thomson Reuters
195 Broadway
New York, NY 10007
US
Phone: 646-822-3000
Email: spencer.giacalone@thomsonreuters.com

Malcolm Betts
ZTE
Email: malcolm.betts@zte.com.cn

Qilei Wang
ZTE
Email: wang.qilei@zte.com.cn

John Drake
Juniper Networks
Email: jdrake@juniper.net

