Internet Engineering Task Force (IETF)
Request for Comments: 7190
Category: Informational
ISSN: 2070-1721

C. Villamizar
Outer Cape Cod Network Consulting

March 2014

Use of Multipath with MPLS and MPLS Transport Profile (MPLS-TP)

Abstract

Many MPLS implementations have supported multipath techniques, and many MPLS deployments have used multipath techniques, particularly in very high-bandwidth applications, such as provider IP/MPLS core networks. MPLS Transport Profile (MPLS-TP) has strongly discouraged the use of multipath techniques. Some degradation of MPLS-TP Operations, Administration, and Maintenance (OAM) performance cannot be avoided when operating over many types of multipath implementations.

Using MPLS Entropy Labels (RFC 6790), MPLS Label Switched Paths (LSPs) can be carried over multipath links while also providing a fully MPLS-TP-compliant server layer for MPLS-TP LSPs. This document describes the means of supporting MPLS as a server layer for MPLS-TP. The use of MPLS-TP LSPs as a server layer for MPLS LSPs is also discussed.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7190.

Table of Contents

1. Introduction

Today the requirement to handle large aggregations of traffic can be met by a number of techniques that we will collectively call "multipath". Multipath applied to parallel links between the same set of nodes includes Ethernet Link Aggregation [IEEE-802.1AX], link bundling [RFC4201], or other aggregation techniques, some of which could be vendor specific. Multipath applied to diverse paths rather than parallel links includes Equal-Cost Multipath (ECMP) as applied to OSPF, IS-IS, or BGP, and equal-cost Label Switched Paths (LSPs). Some vendors support load splitting across equal-cost MPLS LSPs where the load is split proportionally to the reserved bandwidth of the set of LSPs.

RFC 5654 requirement 33 requires the capability to carry a client MPLS Transport Profile (MPLS-TP) or MPLS layer over a server MPLS-TP or MPLS layer [RFC5654]. This is possible in all cases with one exception. When an MPLS LSP exceeds the capacity of any single

component link, it MAY be carried by a network using multipath techniques, but it SHOULD NOT be carried by a single MPLS-TP LSP due to the inherent MPLS-TP capacity limitation imposed by MPLS-TP Operations, Administration, and Maintenance (OAM) fate-sharing constraints and MPLS-TP Loss Measurement OAM packet-ordering constraints (see Section 3.1). Instead, multiple MPLS-TP LSPs SHOULD be used to carry a large MPLS LSP (see Section 4).

The term "composite link" is more general than terms such as "link aggregation" (which is specific to Ethernet) or "ECMP" (which implies equal-cost paths within a routing protocol). The use of the term "composite link" here is consistent with the broad definition in [ITU-T.G.800]. Multipath is very similar to composite link as defined by ITU-T but specifically excludes inverse multiplexing.

MPLS LSPs today are able to function as a server layer and carry client MPLS LSPs. When control-plane signaling is used, forwarding adjacency (FA) advertisements are used to inform the set of Label Switching Routers (LSRs) of Packet Switching Capable (PSC) LSPs within the MPLS topology [RFC4206]. Client MPLS LSP at a higher layer (lower PSC number) may signal their intention to use PSC LSPs as hops in the RSVP-TE Explicit Route Object (ERO). LSRs with no explicit support for RFC 4206 see the PSC LSPs as ordinary links and therefore use them.

An MPLS LSP that has been set up using RSVP-TE appears to its ingress LSR as a viable IP next hop to a distant LSR. If LDP is used and bidirectional RSVP-TE LSP connectivity is available, then LDP signaling can be set up among the RSVP-TE LSP endpoints, and LDP can make use of the RSVP-TE LSP as an LDP hop. This is another form of existing MPLS-in-MPLS use. MPLS LSPs may also make use of hierarchy that is configured through the management plane rather than signaled using RSVP-TE.

These existing forms of MPLS-in-MPLS may traverse multipath hops such as Ethernet Link Aggregation Group (LAG) [IEEE-802.1AX] or MPLS Link Bundling [RFC4201]. MPLS-TP brings with it a new set of requirements not considered in past deployments of the various forms of MPLS-in-MPLS where multipath was in use. This document merely discusses use of existing forwarding and protocol mechanisms that can support the case where either the client-layer LSPs or the server-layer LSPs are MPLS-TP and where multipath is used.

2. Definitions

Please refer to the terminology related to multipath introduced in [ADV-MULTIPATH-REQ]. The following additional terms are used in this document; related terms are grouped together.

Link Bundle
: Link bundling is a multipath technique specific to MPLS [RFC4201]. Link bundling supports two modes of operations. Either an LSP can be placed on one component link of a link bundle, or an LSP can be load-split across all members of the bundle. There is no signaling defined that allows a per-LSP preference regarding load split, therefore whether to load split is generally configured per bundle and applied to all LSPs across the bundle.

All-Ones Component
: Within the context of link bundling, [RFC4201] defines a special case where the same label is to be valid across all component links. This case is indicated in signaling by a bit value of "all ones" when identifying a component link. Following the publication of RFC 4201, for brevity this special case has been referred to as the "all-ones component".

Equal-Cost Multipath (ECMP)
: Equal-Cost Multipath (ECMP) is a specific form of multipath in which the costs of the links or paths must be equal in a given routing protocol. The load may be split equally across all available links (or available paths), or the load may be split proportionally to the capacity of each link (or path).

Loop-Free Alternate Paths (LFA)
: "Loop-free alternate paths" (LFA) are defined in Section 5.2 of RFC 5714 [RFC5714] as follows: "Such a path exists when a direct neighbor of the router adjacent to the failure has a path to the destination that can be guaranteed not to traverse the failure." Further detail can be found in [RFC5286]. LFA as defined for IP Fast Reroute (IPFRR) can be used to load balance by relaxing the equal-cost criteria of ECMP, though IPFRR defined LFA for use in selecting protection paths. When used with IP, proportional split is generally not used. LFA use in load balancing is implemented by some vendors, though it may be rare or non-existent in deployments.

Link Aggregation
: The term "link aggregation" generally refers to Ethernet Link Aggregation as defined by [IEEE-802.1AX]. Ethernet Link Aggregation defines a Link Aggregation Control Protocol (LACP) which coordinates inclusion of Link Aggregation Group (LAG) members in the LAG.

Link Aggregation Group (LAG)
: A group of physical Ethernet interfaces that are treated as a logical link when using Ethernet Link Aggregation is referred to as a Link Aggregation Group (LAG).

LAG Member
: Ethernet Link Aggregation as defined in [IEEE-802.1AX] refers to an individual link in a LAG as a LAG member. A LAG member is a component link. An Ethernet LAG is a composite link. IEEE does not use the terms "composite link" or "component link".

A small set of requirements are discussed. These requirements make use of keywords such as MUST and SHOULD as described in [RFC2119].

3. MPLS as a Server Layer for MPLS-TP

An MPLS LSP may be used as a server layer for MPLS-TP LSPs as long as all MPLS-TP requirements are met. Section 3.1 reviews the basis for requirements of a server layer that supports MPLS-TP as a client layer. Key requirements include OAM "fate-sharing" and that packets within an MPLS-TP LSP (including both payload and OAM packets) not be reordered. Section 3.2 discusses implied requirements where MPLS is the server layer for MPLS-TP client LSPs and describes a set of solutions that use existing MPLS mechanisms.

3.1. MPLS-TP Forwarding and Server-Layer Requirements

[RFC5960] defines the data-plane requirements for MPLS-TP. Two very relevant paragraphs in Section 3.1.1 ("LSP Packet Encapsulation and Forwarding") are the following:

RFC 5960, Section 3.1.1, Paragraph 3
: Except for transient packet reordering that may occur, for example, during fault conditions, packets are delivered in order on L-LSPs, and on E-LSPs within a specific ordered aggregate.

RFC 5960, Section 3.1.1, Paragraph 6

> Equal-Cost Multi-Path (ECMP) load-balancing MUST NOT be performed on an MPLS-TP LSP. MPLS-TP LSPs as defined in this document MAY operate over a server layer that supports load-balancing, but this load-balancing MUST operate in such a manner that it is transparent to MPLS-TP. This does not preclude the future definition of new MPLS-TP LSP types that have different requirements regarding the use of ECMP in the server layer.

[RFC5960], Section 3.1.1, Paragraph 3 requires that packets within a specific ordered aggregate be delivered in order. This same requirement is already specified by Differentiated Services [RFC2475]. [RFC5960], Section 3.1.1, Paragraph 6 explicitly allows a server layer to use ECMP, provided that it is transparent to the MPLS-TP client layer.

[RFC6371] adds a requirement for data traffic and OAM traffic "fate-sharing". The following paragraph in Section 1 ("Introduction") summarizes this requirement.

RFC 6371, Section 1, Paragraph 7

> OAM packets that instrument a particular direction of a transport path are subject to the same forwarding treatment (i.e., fate-share) as the user data packets and in some cases, where Explicitly TC-encoded-PSC LSPs (E-LSPs) are employed, may be required to have common per-hop behavior (PHB) Scheduling Class (PSC) End-to-End (E2E) with the class of traffic monitored. In case of Label-Only-Inferred-PSC LSP (L-LSP), only one class of traffic needs to be monitored, and therefore the OAM packets have common PSC with the monitored traffic class.

[RFC6371] does not prohibit multilink techniques in Section 4.6 ("Fate-Sharing Considerations for Multilink"), where multilink is defined as Ethernet Link Aggregation and the use of Link Bundling for MPLS, but it does declare that such a network would be only partially MPLS-TP compliant. The characteristic that is to be avoided is contained in the following sentence in that section.

RFC 6371, Section 4.6, Paragraph 1, last sentence

> These techniques frequently share the characteristic that an LSP may be spread over a set of component links and therefore be reordered, but no flow within the LSP is reordered (except when very infrequent and minimally disruptive load rebalancing occurs).

A declaration that implies that Link Bundling for MPLS yields a partially MPLS-TP-compliant network is perhaps overstated since only the Link Bundling all-ones component link has this characteristic.

[RFC6374] defines a direct Loss Measurement (LM) where LM OAM packets cannot be reordered with respect to payload packets. This will require that payload packets themselves not be reordered. The following paragraph in Section 2.9.4 ("Equal Cost Multipath") gives the reason for this restriction.

RFC 6374, Section 2.9.4, Paragraph 2
> The effects of ECMP on loss measurement will depend on the LM mode. In the case of direct LM, the measurement will account for any packets lost between the sender and the receiver, regardless of how many paths exist between them. However, the presence of ECMP increases the likelihood of misordering both of LM messages relative to data packets and of the LM messages themselves. Such misorderings tend to create unmeasurable intervals and thus degrade the accuracy of loss measurement. The effects of ECMP are similar for inferred LM, with the additional caveat that, unless the test packets are specially constructed so as to probe all available paths, the loss characteristics of one or more of the alternate paths cannot be accounted for.

3.2. Methods of Supporting MPLS-TP Client LSPs over MPLS

Supporting MPLS-TP LSPs over a fully MPLS-TP conformant MPLS LSP server layer where the MPLS LSPs are making use of multipath requires special treatment of the MPLS-TP LSPs such that those LSPs meet MPLS-TP forwarding requirements (see Section 3.1). This implies the following brief set of requirements.

MP#1 It MUST be possible for a midpoint MPLS-TP Label Switching Router (LSR) that is serving as ingress to a server-layer MPLS LSP to identify MPLS-TP LSPs, so that MPLS-TP forwarding requirements can be applied, or to otherwise accommodate the MPLS-TP forwarding requirements.

MP#2 The ability to completely exclude MPLS-TP LSPs from the multipath hash and load split SHOULD be supported. If the selected component link no longer meets requirements, an LSP is considered down, which may trigger protection and/or may require that the ingress LSR select a new path and signal a new LSP.

MP#3 It SHOULD be possible to ensure that MPLS-TP LSPs will not be moved to another component link as a result of a load-rebalancing operation for multipath. If the selected component link no longer meets requirements, another component link may be selected; however, a change in path SHOULD NOT occur solely for load balancing.

MP#4 Where a Resource Reservation Protocol - Traffic Engineering (RSVP-TE) control plane is used, it MUST be possible for an ingress LSR that is setting up an MPLS-TP or an MPLS LSP to determine at path selection time whether a link or Forwarding Adjacency (FA; see [RFC4206]) within the topology can support the MPLS-TP requirements of the LSP.

The reason for requirement MP#1 may not be obvious. An MPLS-TP LSP may be aggregated along with other client LSPs by a midpoint LSR into a very large MPLS server-layer LSP, as would be the case in a core-node-to-core-node MPLS LSP between major cities. In this case, the ingress of the MPLS LSP, being a midpoint LSR for a set of client LSPs, has no signaling mechanism that can be used to determine whether one of its specific client LSPs is using MPLS or MPLS-TP. Multipath load splitting can be avoided for MPLS-TP LSPs if at the MPLS server-layer LSP ingress LSR an Entropy Label Indicator (ELI) and Entropy Label (EL) are added to the label stack by the midpoint LSR for the client MPLS-TP LSP, at the ingress of the MPLS LSP [RFC6790]. For those client LSPs that are MPLS-TP LSPs, a single per-LSP EL value must be chosen. For those client LSPs that are MPLS LSPs, per-packet entropy below the top label must, for practical reasons, be used to determine the entropy label value. The resulting label stack contains the server MPLS LSP label, ELI, EL and the client LSP label. Requirement MP#1 simply states that there must be a means to make this decision.

There is currently no signaling mechanism defined to support requirement MP#1, though that does not preclude a new extension being defined later. In the absence of a signaling extension, MPLS-TP can be identified through some form of configuration, such as configuration that provides an MPLS-TP-compatible server layer to all LSPs arriving on a specific interface or originating from a specific set of ingress LSRs.

Alternatively, the need for requirement MP#1 can be eliminated if every MPLS-TP LSP created by an MPLS-TP ingress makes use of an Entropy Label Indicator (ELI) and Entropy Label (EL) below the MPLS-TP label [RFC6790]. This would require that all MPLS-TP LSRs in a deployment support Entropy Label, which may render it impractical in many deployments.

Some hardware that exists today can support requirement MP#2. Signaling in the absence of MPLS Entropy Labels can make use of link bundling with the path pinned to a specific component for MPLS-TP LSPs and link bundling using the all-ones component for MPLS LSPs. This prevents MPLS-TP LSPs from being carried within MPLS LSPs but does allow the coexistence of MPLS-TP and very large MPLS LSPs.

When Entropy Label Indicators (ELIs) and Entropy Labels (ELs) are not applied by MPLS-TP ingresses, MPLS-TP LSPs can be carried as client LSPs within an MPLS server LSP if the ingress of the MPLS server-layer LSP pushes an Entropy Label Indicator (ELI) and Entropy Label (EL) below the server-layer LSP label(s) in the label stack, just above the MPLS-TP LSP label entry [RFC6790]. The value of EL can be randomly selected at the client MPLS-TP LSP setup time, and the same EL value can be used for all packets of that MPLS-TP LSP. This allows MPLS-TP LSPs to be carried as client LSPs within MPLS LSPs and satisfies MPLS-TP forwarding requirements but requires that MPLS LSRs be able to identify MPLS-TP LSPs (requirement MP#1).

MPLS-TP traffic can be protected from degraded performance due to an imperfect load split if the MPLS-TP traffic is given queuing priority. For example, using (1) strict priority and policing, shaping at ingress, or per-LSP shaping locally, or (2) per-LSP weighted queuing locally. This can be accomplished using the Traffic Class (TC) field and Diffserv treatment of traffic [RFC5462] [RFC2475]. In the event of congestion due to load imbalance, only non-prioritized traffic will suffer as long as there is a low percentage of prioritized traffic.

If MPLS-TP LSPs are carried within MPLS LSPs and ELI and EL are used, requirement MP#3 is satisfied (1) for uncongested links where load balancing is not required, or (2) for MPLS-TP LSPs using Traffic Class (TC) and Diffserv, where the load rebalancing implementation rebalances only the less preferred traffic. Load rebalance is generally needed only when congestion occurs; therefore, restricting MPLS-TP to be carried over MPLS LSPs that are known to traverse only links that are expected to be uncongested can satisfy requirement MP#3.

An MPLS-TP LSP can be pinned to a Link Bundle component link if the behavior of requirement MP#2 is preferred. An MPLS-TP LSP can be assigned to a Link Bundle but not pinned if the behavior of requirement MP#3 is preferred. In both of these cases, the MPLS-TP LSP must be the top-level LSP, except as noted above.

If MPLS-TP LSPs can be moved among component links, then the Link Bundle all-ones component link can be used or server-layer MPLS LSPs can be used with no restrictions on the server-layer MPLS use of multipath, except that Entropy Labels must be supported along the entire path. An Entropy Label must be used to ensure that all of the MPLS-TP payload and OAM traffic are carried on the same component, except during very infrequent transitions due to load balancing. Since the Entropy Label Indicator and Entropy Label are always placed above the Generic Associated Channel Label (GAL) in the stack, the

presence of a GAL will not affect the selection of a component link as long as the LSR does not hash on the label stack entries below the Entropy Label.

An MPLS-TP LSP may not traverse multipath links on the path where MPLS-TP forwarding requirements cannot be met. Such links include any using pre-[RFC6790] Ethernet Link Aggregation, pre-[RFC6790] Link Bundling using the all-ones component link, or any other form of multipath that does not support termination of the entropy search at the EL as called for in [RFC6790]. An MPLS-TP LSP MUST NOT traverse a server-layer MPLS LSP that traverses any form of multipath that does not support termination of the entropy search at the EL. For this to occur, the MPLS-TP ingress LSR MUST be aware of these links. This is the reason for requirement MP#4.

Requirement MP#4 can be supported using administrative attributes. Administrative attributes are defined in [RFC3209]. Some configuration is required to support this.

In MPLS Link Bundling the requirement for bidirectional co-routing can be interpreted as meaning that the same set of LSRs must be traversed or can be interpreted to mean that the same set of component links must be traversed [RFC4201] [RFC3473]. Following the procedures of Section 3 of RFC 3473 where Link Bundling is used only ensures that the same set of LSRs are traversed and that acceptable labels are created in each direction.

When an MPLS-TP LSP is set up over a MPLS LSP, if the MPLS-TP LSP is a bidirectional LSP, then providers who want to only set these MPLS-TP LSPs over bidirectional co-routed MPLS LSPs can make use of administrative attributes [RFC3209] to ensure that this occurs. If MPLS-TP LSPs are carried by unidirectional MPLS LSPs, the MPLS-TP OAM will be unaffected, as only the MPLS LSP endpoints will appear as MPLS-TP OAM Maintenance Entity Group Intermediate Points (MIPs).

Two methods of adding an Entropy Label are described above. The MPLS-TP ingress must have a means to determine which links can support MPLS-TP in selecting a path (MP#4). Administrative attributes can satisfy that requirement. If the MPLS-TP LSR is capable of adding ELI/EL to the label stack, this method is preferred. However, equipment furthest from a provider’s network core is the least likely to support RFC 6790 in the near term. For portions of the topology where an MPLS-TP is carried within a server-layer MPLS LSP, the ingress of the server-layer MPLS LSP can add ELI/EL using a fixed EL value per client LSP, except those known not to require MPLS-TP treatment. There are numerous ways to determine which client LSPs are MPLS-TP LSPs and which are not. While this

determination is out of scope and will vary among deployments, configuration or the presence of specific attribute affinities in RSVP-TE signaling are among the likely means to do so.

4. MPLS-TP as a Server Layer for MPLS

Carrying MPLS LSPs that are larger than a component link over an MPLS-TP server layer requires that the large MPLS client-layer LSP be accommodated by multiple MPLS-TP server-layer LSPs. MPLS multipath can be used in the client-layer MPLS.

Creating multiple MPLS-TP server-layer LSPs places a greater Incoming Label Map (ILM) scaling burden on the LSR. High-bandwidth MPLS cores with a smaller amount of nodes have the greatest tendency to require LSPs in excess of component links; therefore, the reduction in the number of nodes offsets the impact of increasing the number of server-layer LSPs in parallel. Today, only in cases where deployed LSR ILMs are small would this be an issue.

The most significant disadvantage of MPLS-TP as a server layer for MPLS is that the use of MPLS-TP server-layer LSPs reduces the efficiency of carrying the MPLS client layer. The service that provides by far the largest offered load in provider networks is the Internet, for which the LSP capacity reservations are predictions of expected load. Many of these MPLS LSPs may be smaller than component link capacity. Using MPLS-TP as a server layer results in bin-packing problems for these smaller LSPs. For those LSPs that are larger than component link capacity, the LSP capacities need not be (and often are not) integer multiples of convenient capacity increments such as 10 Gbit/s. Using MPLS-TP as an underlying server layer greatly reduces the ability of the client-layer MPLS LSPs to share capacity. For example, when one MPLS LSP is underutilizing its predicted capacity, the fixed allocation of MPLS-TP to component links may not allow another LSP to exceed its predicted capacity. Using MPLS-TP as a server layer may result in less efficient use of resources and may result in a less cost-effective network.

No additional requirements beyond MPLS-TP as it is now currently defined are required to support MPLS-TP as a server layer for MPLS. It is therefore viable but has some undesirable characteristics discussed above.

5. Summary

MPLS equipment deployed in the core currently supports multipath. For large service providers, core LSR must support some form of multipath to be deployable. Deployed MPLS access and edge equipment is often oblivious to the use of multipath in the core. It is

expected that at least first-generation MPLS-TP equipment will be oblivious to the use of multipath in the core. This first-generation MPLS-TP equipment is deployable in a core using multipath, with no adverse impact to RSVP-TE signaling, if:

1. the edge equipment can support administrative attributes (RFC 3209),

2. the core equipment can support ELI/EL, and

3. the core equipment can put a per-LSP fixed EL value on any LSP that indicates a particular attribute affinity or can identify a client MPLS-TP LSP through some other means.

There are no issues carrying MPLS over MPLS-TP, except when the MPLS LSP is too big to be carried by a single MPLS-TP LSP. Most MPLS core equipment and some edge equipment can configure an MPLS Link Bundle [RFC4201] over multiple component links where the component links are themselves MPLS LSP. This existing capability can be used to carry large MPLS LSPs and overcome the limited capacity of any single server-layer MPLS-TP LSP.

MPLS OAM and MPLS-TP OAM are unaffected in the following cases proposed in this document:

1. Where MPLS is carried over a single MPLS-TP, all traffic flows on one link, MPLS OAM is unaffected and need not use multipath support in LSP Ping [RFC4379].

2. Where MPLS-TP is carried over MPLS, all traffic for that MPLS-TP LSP is carried over one link thanks to the fixed EL value. In this case, MPLS-TP OAM is unaffected.

3. Where MPLS LSPs are carried over MPLS LSPs (an existing case) or over multiple MPLS-TP LSPs, the multipath support in LSP Ping is used and LSP Ping operation is unaffected [RFC4379] [RFC6425].

6. Acknowledgements

Carlos Pignataro, Dave Allan, and Mach Chen provided valuable comments and suggestions. Carlos suggested that MPLS-TP requirements in RFC 5960 be explicitly referenced or quoted. An email conversation with Dave led to the inclusion of references and quotes from RFCs 6371 and 6374. Mach made suggestions to improve the clarity of the document.

7. Security Considerations

This document specifies use of existing MPLS and MPLS-TP mechanisms to support MPLS and MPLS-TP as client and server layers for each other. This use of existing mechanisms supports coexistence of MPLS/GMPLS (without MPLS-TP) when used over a packet network, MPLS-TP, and multipath. The combination of MPLS, MPLS-TP, and multipath does not introduce any new security threats. The security considerations for MPLS/GMPLS and for MPLS-TP are documented in [RFC5920] and [RFC6941].

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.

[RFC5654] Niven-Jenkins, B., Brungard, D., Betts, M., Sprecher, N., and S. Ueno, "Requirements of an MPLS Transport Profile", RFC 5654, September 2009.

[RFC5960] Frost, D., Bryant, S., and M. Bocci, "MPLS Transport Profile Data Plane Architecture", RFC 5960, August 2010.

[RFC6371] Busi, I. and D. Allan, "Operations, Administration, and Maintenance Framework for MPLS-Based Transport Networks", RFC 6371, September 2011.

[RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", RFC 6374, September 2011.

[RFC6790] Kompella, K., Drake, J., Amante, S., Henderickx, W., and L. Yong, "The Use of Entropy Labels in MPLS Forwarding", RFC 6790, November 2012.

8.2. Informative References

[ADV-MULTIPATH-REQ]
Villamizar, C., McDysan, D., Ning, S., Malis, A., and L. Yong, "Requirements for Advanced Multipath in MPLS Networks", Work in Progress, February 2014.

[IEEE-802.1AX]
IEEE, "IEEE Standard for Local and Metropolitan Area Networks - Link Aggregation", IEEE Std 802.1AX-2008, 2006, <http://standards.ieee.org/getieee802/download/802.1AX-2008.pdf>.

[ITU-T.G.800]
ITU-T, "Unified functional architecture of transport networks", ITU-T G.800, 2007, <http://www.itu.int/rec/T-REC-G/recommendation.asp?parent=T-REC-G.800>.

[RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and W. Weiss, "An Architecture for Differentiated Services", RFC 2475, December 1998.

[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.

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

[RFC4201] Kompella, K., Rekhter, Y., and L. Berger, "Link Bundling in MPLS Traffic Engineering (TE)", RFC 4201, October 2005.

[RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", RFC 4206, October 2005.

[RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.

[RFC5286] Atlas, A. and A. Zinin, "Basic Specification for IP Fast Reroute: Loop-Free Alternates", RFC 5286, September 2008.

[RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", RFC 5462, February 2009.

[RFC5714] Shand, M. and S. Bryant, "IP Fast Reroute Framework", RFC 5714, January 2010.

[RFC5920] Fang, L., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

[RFC6425] Saxena, S., Swallow, G., Ali, Z., Farrel, A., Yasukawa, S., and T. Nadeau, "Detecting Data-Plane Failures in Point-to-Multipoint MPLS - Extensions to LSP Ping", RFC 6425, November 2011.

[RFC6941] Fang, L., Niven-Jenkins, B., Mansfield, S., and R. Graveman, "MPLS Transport Profile (MPLS-TP) Security Framework", RFC 6941, April 2013.

Author’s Address

Curtis Villamizar
Outer Cape Cod Network Consulting

EMail: curtis@occnc.com