

Routing Working Group
Internet Draft
Intended Status: Standard
Expires:

N. So
A. Malis
D. McDysan
Verizon
L. Yong
Huawei
F. Jounay
France Telecom
Y. Kamite
NTT
October 16, 2009

Requirements for MPLS Over Composite Link

[draft-so-yong-mpls-ctg-requirement-00](#)

Status of this Memo

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

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). This document may not be modified, and derivative works of it may not be created, and it may not be published except as an Internet-Draft.

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). This document may not be modified, and derivative works of it may not be created, except to publish it as an RFC and to translate it into languages other than English.

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 December 17, 2009.

Copyright Notice

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Abstract

This document presents motivation, a problem statement and operational requirements for a traffic distribution problem in today's IP/MPLS network when multiple links are configured between two routers. It defines a composite link as a group of parallel links that can be considered as a single traffic engineering link or as an IP link, and used for MPLS. The framework described in a companion document is used to organize the requirements.

Table of Contents

1.	Introduction.....	3
2.	Conventions used in this document	4
2.1.	Acronyms	4
2.2.	Terminology.....	4
3.	Motivation and Summary Problem Statement.....	5
3.1.	Motivation.....	5
3.2.	Summary of Problems Requiring Solution	7
4.	Requirements.....	7
4.1.	Interior Functions	8
4.1.1.	Functions common to all LSP flows	8
4.1.1.1.	Traffic Flow and Connection Mapping.....	9
4.1.1.2.	Management of Other Operational Aspects.....	9
4.1.1.2.1.	Resilience.....	9
4.1.1.2.2.	Fault management requirement	9
4.1.1.2.3.	Flow/Connection Mapping Change Frequency.	10
4.1.2.	Functions specific to LSP flows with TE information	11
4.1.3.	Functions specific to LSP flows without TE information	12
4.1.4.	Sets of LSP flows with and without TE information..	12
4.1.4.1.	Handling Bandwidth Shortage Events.....	13
4.2.	Exterior Functions	13

4.2.1. Functions common to all LSP flows	13
4.2.1.1. Signaling Protocol Extensions.....	13
4.2.1.2. Routing Advertisement Extensions.....	13
4.2.2. Functions specific to LSP flows with TE information	14
4.2.2.1. Signaling Protocol Extensions Requirements...	14
4.2.2.2. Routing Advertisement Extensions.....	15
4.2.3. Functions specific to LSP flows without TE information	15
4.2.3.1. Signaling Protocol Extensions.....	15
4.2.3.2. Routing Advertisement Extensions.....	16
4.2.4. Functions specific to LSP flows with and without TE information.....	16
4.2.4.1. Signaling Protocol Extensions.....	16
4.2.4.2. Routing Advertisement Extensions.....	16

5. Security Considerations	16
6. IANA Considerations.....	16
7. References.....	17
7.1. Normative References.....	17
7.2. Informative References.....	17
8. Acknowledgments.....	18

[1](#). Introduction

IP/MPLS network traffic growth forces carriers to deploy multiple parallel physical/logical links between adjacent routers as the total capacity of all aggregated traffic flows exceed the capacity of a single link. The network is expected to carry aggregated traffic flows some of which approach the capacity of any single link, and also some flows that may be very small compared to the capacity of a single link.

Operating an MPLS network with multiple parallel links between all adjacent routers causes scaling problems in the routing protocols. This issue is addressed in [[RFC4201](#)] which defines the notion of a Link Bundle -- a set of identical parallel traffic engineered (TE) links (called component links) that are grouped together and advertised as a single TE link within the routing protocol.

The Link Bundle concept is somewhat limited because of the requirement that all component links must have identical capabilities, and because it applies only to TE links. This document sets out a more generic set of requirements for grouping together a set of parallel data links that may have different characteristics, and for advertising and operating them as a single TE or non-TE link called a Composite Link.

First, there is a set of requirements related to the interior functioning of a router, of which the operational requirement is to have consistent configuration, reporting and alarm notification. The functions that impact the routers other than those connected by the composite link are control plane routing and signaling protocols. A further subdivision of the requirements is based upon characteristics of the combination of MPLS signaling protocols in use, namely RSVP-TE only, LDP only, or a combination of RSVP-TE and LDP. Extension requirements to IGP-related protocols are also described in this document. Furthermore, there are requirements that relate to use of signaling (and possibly routing) that can be used within the same layer or between layers.

Applicability of the work within this document is focused on MPLS traffic as controlled through control plane protocols. Thus, this document describes requirements related to the routing protocols that advertise link parameters and the Resource Reservation Protocol (RSVP-TE) and the Label Distribution Protocol (LDP) signaling protocols that distribute MPLS labels and establish Label Switched Paths (LSPs). Interactions between the control plane and the data and

management planes are also addressed. The focus of this document is on MPLS traffic either signaled by RSVP-TE or LDP. IP traffic over multiple parallel links is handled relatively well by ECMP or LAG/hashing methods. The handling of IP control plane traffic is within the scope of the requirements of this document.

A companion framework document [[CLFRAMEWORK](#)] further describes the overall context, a structure, and some standardization considerations. Specific protocol solutions are outside the scope of this document.

[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](#).

[2.1.](#) Acronyms

BW: Bandwidth

ECMP: Equal Cost Multi-Path

FRR: Fast Re-Route

LAG: Link Aggregation Group

LDP: Label Distribution Protocol

LSP: Label Switched Path

MPLS: Multi-Protocol Label Switching

OAM: Operation, Administration, and Management

PDU: Protocol Data Unit

PE: Provider Edge device

RSVP: ReSource reSeRvAtion Protocol

RTD: Real Time Delay

TE: Traffic Engineering

VRF: Virtual Routing and Forwarding

[2.2.](#) Terminology

Composite Link: a group of component links, which can be considered as a single MPLS TE link or as a single IP link used for MPLS.

Component Link: a physical link (e.g., Lambda, Ethernet PHY, SONET/SDH, OTN, etc.) with packet transport capability, or a logical link (e.g., MPLS LSP, Ethernet VLAN, MPLS-TP LSP, etc.)

Connection: An aggregation of traffic flows which are treated together as a single unit by the CTG Interior Function for the purpose of routing onto a specific component link and measuring traffic volume.

Exterior Functions: These are performed by an MPLS router that makes a composite link useable by the network via control protocols, or by an MPLS router that interacts with other routers to dynamically control a component link as part a composite link. These functions are those that interact via routing and/or signaling protocols with other routers in the same layer network or other layer networks.

Interior Functions: Actions performed by the MPLS routers directly connected by a composite link. This includes the determination of the connection and component link on which a traffic flow is placed.

Although a local implementation matter, the configuration control of certain aspects of these interior functions is an important operational requirement.

Traffic Flow: A set of packets that with common identifier characteristics that the composite link is able to use to aggregate traffic into Connections. Identifiers can be an MPLS label stack or any combination of IP addresses and protocol types for routing, signaling and management packets.

Virtual Interface: Composite link characteristics advertised in IGP

[3.](#) Motivation and Summary Problem Statement

[3.1.](#) Motivation

There are several established approaches to using multiple parallel links between a pair of routers. These have limitations as summarized below.

- o ECMP/Hashing/LAG: IP traffic composed of a large number of flows with bandwidth that is small with respect to the individual link capacity can be handled relatively well using ECMP/LAG approaches. However, these approaches do not make use of MPLS control plane information nor traffic volume information. Distribution techniques applied only within the data plane can result in less than ideal load balancing across component links of a composite link.

-
- o Advertisement of each component link into the IGP. Although this would address the problem, it has a scaling impact on IGP routing, and was an important motivation for the specification of link bundling [[RFC4201](#)]. However, link bundling does not support a set of component links with different characteristics (e.g., bandwidth, latency) and only supports RSVP-TE.
 - o Inverse Multiplexing: Making multiple parallel links to appear as a single link using inverse multiplexing techniques, such as proposals under discussion in the [[PWBONDING](#)]. However, the inverse multiplexed link will have a latency of the link with the

largest latency. When there is a mix of latency sensitive traffic with other traffic that is less sensitive to latency, having all traffic experience the latency of the worst link is not acceptable to operators.

- o Planning Tool LSP Assignment: Although theoretically optimal, an external system that participates in the IGP, measures traffic and assigns TE LSPs and/or adjusts IGP metrics has a potentially large response time to certain failure scenarios. Furthermore, such a system could make use of more information than provided by link bundling IGP advertisements and could make use of mechanisms that would allow pinning MPLS traffic to a particular component link in a CTG.
- o In a multi-layer network, the characteristics of a component link can be altered by a lower layer network and this can create significant operational impact in some cases. For example, if a lower layer network performs restoration and markedly increases the latency of a link in a link bundle, the traffic placed on this longer latency link may generate user complaints and/or exceed the parameters of a Service Level Agreement (SLA).
- o In the case where multiple routing instances could share a composite link, inefficiency can result if either 1) specific component links are assigned to an individual routing instance, or 2) if statically assigned capacity is made to a logical/sub interface in each component link of a composite link for each routing instance. In other words, the issue is that unused capacity in one routing instance cannot be used by another in either of these cases.
- o Unicast LDP signaled LSP flows follow the IGP determined topology in a multipoint-to-point manner. The principal means of control of LDP flows is through adjustment of the IGP metric. The simplicity of this method is attractive. However, this means that the LDP flow traffic level can change significantly in response to some link and/or node failures, thus significantly change the traffic level at points within a network. A means for operators to better manage LDP signaled flows is therefore highly desirable.

3.2. Summary of Problems Requiring Solution

The following bullets highlight aspects of solutions for which detailed requirements are stated in [Section 5](#).

- o Ensure the ability to transport both RSVP-TE and LDP signaled non-TE LSPs on the same composite link (i.e., a single set of component links) while maintaining acceptable service quality for both RSVP-TE and LDP signaled LSPs.
- o Extend a link bundling type function to scenarios with groups of links having different characteristics (e.g., bandwidth, latency).
- o When an end-to-end LSP signaled by RSVP-TE uses a composite link, the composite link must select a component link that meets the end-to-end requirements for the LSP. To perform this function, the network elements at the end of a composite link must be made aware of the required, desired, and acceptable link characteristics (e.g., latency, optimization frequency) for each hop in the path. The solution should be able to operate in a statically configured or dynamically signaled manner.
- o Support sets of component links between routers across intermediate nodes at the same and/or lower layers where the characteristics (e.g., latency) of said links may change dynamically. The solution should support the case where the changes in characteristics of these links are not communicated by the IGP (e.g., a link in a lower layer network has a change in latency or QoS due to a restoration action). The solution could measure the performance (e.g., latency), and/or receive the results of a measurement or computation from lower layer configuration data via signaling.

4. Requirements

As defined in the terminology section a (traffic) flow is the smallest unit of traffic assignable to a connection, and connections are assigned to a component link that is part of a composite link. The composite link has routing information, normal IGP that has no TE information and IGP with TE extensions (IGP-TE) and signaling information with TE information (RSVP-TE) and without TE information. The following table summarizes the three cases covered in this requirements section.

Flows	IGP	IGP-TE	RSVP-TE	LDP
-----	---	-----	-----	---
With TE Info	Y	Y	Y	N
Without TE Info	Y	N	N	Y
With & Without TE Info	Y	Y	Y	Y

Furthermore, if a requirement would be repeated for each of the above three cases (e.g., IGP related routing information) it is described in a section common to all flows.

Therefore, the outline of this section is structured as follows:

- o Management/Measurement of Interior Functions
 - Functions common to all LSP flows
 - Functions specific to LSP flows with TE information
 - Functions specific to LSP flows without TE information
 - Sets of LSP flows with and without TE information
- o Exterior Functions
 - Functions common to all LSP flows
 - Functions specific to LSP flows with TE information
 - Functions specific to LSP flows without TE information
 - Sets of LSP flows with and without TE information

[4.1.](#) Interior Functions

[4.1.1.](#) Functions common to all LSP flows

The operator SHALL be able to configure a "virtual interface" corresponding to a composite link that has at least all of the normal IGP routing parameters .

The solution SHALL allow configuration of virtual interface IGP parameters for a non-TE (i.e., normal IP routing) link used for MPLS (e.g., administrative cost or weight).

- o The solution SHALL support configuration of a composite link

composed of set of component links that may be logical or physical.

The "virtual interface" SHALL appear as a fully-featured routing adjacency not just as an FA [[RFC4206](#)]. In particular, it needs to work with at least the following IP/MPLS control protocols: OSPF/IS-IS, LDP, IGPOSPF-TE/ISIS-TE, and RSVP-TE.

The solution SHALL accept a new component link or remove an existing component link by operator provisioning or in response to signaling at a lower layer (e.g., using GMPLS).

The solution SHALL support derivation of the advertised interface parameters from configured component link parameters based on operator policy.

A composite link SHALL be configurable as a numbered or unnumbered link (virtual interface in IP/MPLS).

A component link SHALL be configurable as a numbered link or unnumbered link. A component link should be not advertised in IGP.

[4.1.1.1](#). Traffic Flow and Connection Mapping

The solution SHALL support operator assignment of traffic flows to specific connections.

The solution SHALL support operator assignment of connections to specific component links.

The solution SHALL support IP packet transport across a composite link for control plane (signaling, routing) and management plane functions.

In order to prevent packet loss, the solution must employ make-before-break when a change in the mapping of a connection to a component link mapping change has to occur.

[4.1.1.2](#). Management of Other Operational Aspects

[4.1.1.2.1](#). Resilience

The component link recovery scheme SHALL perform equal to or better than existing local recovery methods. A short service disruption may occur during the recovery period. Fast ReRoute (FRR) SHALL be

configurable for a composite link.

As part of FRR, a method to recover from CTG node failure is desirable. OAM Messaging Support

[4.1.1.2.2](#). Fault management requirement

There are two aspects of fault management in the solution. One is about composite link between two local adjacent routers. The other is about the individual component link.

So, et al.

Expires April 16, 2010

[Page 9]

OAM protocols for fault management from the outside routers (e.g., LSP-Ping/Trace, IP-ping/Trace) SHALL be transparently treated.

For example, it is expected that LSP-ping/trace message is able to diagnose composite link status and its associated virtual interface information; however, it is not required to directly treat individual component link and CTG-connection because they are local matter of two routers.

The solution SHALL support fault notification mechanism (e.g., syslog, SNMP trap to the management system/operators) with the granularity level of affected part as detailed below:

- o Data-plane of component link level
- o Data-plane of composite link level (as a whole, and as a sum of associated component links)
- o Control-plane of the virtual interface level (i.e., routing/signaling on it)
- o A composite link that believes that the underlying component link server layers might not efficiently report failures, can run Bidirectional Forwarding Detection (BFD) at the component link layer.

Race conditions: The solution shall support configuration of timers so that lower layer methods have time to detect/restore faults before a composite link function would be invoked.

The solution SHALL allow operator or control plane to query the component link to which an LSP is assigned.

[4.1.1.2.3](#). Flow/Connection Mapping Change Frequency

The solution requires methods to dampen the frequency of flow to connection mapping change, connection bandwidth change, and/or connection to component link mapping changes (e.g., for re-optimization). Operator imposed control policy regarding the frequency of change for sets of flows SHALL be supported.

The solution SHALL support latency and delay variation sensitive traffic and limit the mapping change for these flows, and place them on component links that have lower latency.

The determination of latency sensitive traffic SHALL be determined by any of the following methods:

- o Use of a pre-defined local policy setting at composite link ingress that can be mapped to a flow

- o A manually configured setting at composite link ingress that can be mapped to a flow

The determination of latency sensitive traffic SHOULD be determined (if possible) by any of the following methods:

- o Pre-set bits in the Payload (e.g., DSCP bits for IP or Ethernet user priority for Ethernet payload) which are typically assigned by end-user
- o MPLS Traffic-Class Field (aka EXP) which is typically mapped by the LER/LSR on the basis that its value is given for differentiating latency-sensitive traffic of end-users

[4.1.2.](#) Functions specific to LSP flows with TE information

The following requirements apply to the case of RSVP-TE signaled LSPs where the virtual interface is configured with IGP TE extensions.

The solution SHALL allow configuration of virtual interface parameters for TE extensions link (e.g., available bandwidth, maximum bandwidth, maximum allowable LSP bandwidth, TE metric, and resource classes (i.e., administrative groups) or link colors).

The solution SHALL support configuration of a composite link composed of set of component links , with each component link potentially having at least the following characteristics which may differ:

- o Bandwidth
- o Latency
- o QoS characteristics (e.g., jitter, error rate)

The solution SHALL support the admission control by RSVP-TE that is signaled from the routers outside the composite link. Note that RSVP-TE signaling need not specify the actual component link to be used? because the selection of component link is the local matter of two adjacent routers based upon signaled and locally configured information.

The solution shall be able to receive, interpret and act upon at least the following RSVP-TE signaled parameters: bandwidth setup priority, and holding priority [RFC 3209, [RFC 2215](#)] preemption priority and traffic class [[RFC 4124](#)], and apply them to the CTG connections where the LSP is mapped.

The solution shall support configuration of at least the following parameters on a per composite link basis:

- o Local Bandwidth Oversubscription factor

The determination of latency sensitive traffic SHALL be determined by any of the following methods:

- o MPLS traffic class in a RSVP-TE signaling message (i.e., Diffserv-TE traffic class [[RFC 4124](#)])

[4.1.3](#). Functions specific to LSP flows without TE information

The following requirements apply to the case of LDP signaled LSPs when no signaled TE information is available.

The solution shall map LDP-assigned labeled packets based upon local configuration (e.g., label stack depth) to define a connection that is mapped to one of the component link.

The solution SHALL map LDP-assigned labeled packets that identify the outer label's FEC.

The solution SHALL support entropy labels [Entropy Label] to map more granular flows to connections.

The solution SHALL be able to measure the bandwidth actually used by a particular connection and derive proper local traffic TE information for the connection.

When the connection bandwidth exceeds the component link capacity, the solution SHALL be able to reassign the traffic flows to several connections.

The solution SHALL support management plane controlled parameters that define at least a minimum bandwidth, maximum bandwidth, preemption priority, and holding priority for each connection without TE information (i.e., LDP signaled LSP that does not contain the same information as an RSVP-TE signaled LSP).

4.1.4. Sets of LSP flows with and without TE information

The solution shall support separation of resources for traffic flows mapped to connections that have access to TE information (e.g., RSVP-TE signaled flows) from those that do not have access to TE information (e.g., LDP-signaled flows or RSVP-TE LSP with Zero bandwidth).

Component links in a composite link may fail independently. The failure of a component link may impact some connections. The impacted connections shall be transferred to other active component links using the same rules as for the original assignment of connections to component links. In the event that all connections cannot be recovered, configured priority and preemption parameters SHOULD be employed.

4.1.4.1. Handling Bandwidth Shortage Events

The following requirements apply to a virtual interface that supports traffic flows both with and without TE information, in response to a bandwidth shortage event. A "bandwidth shortage" can arise in a composite link if the total bandwidth of the connections with provisioned/signaled TE information and those signaled without TE information (but with measured bandwidth) exceeds the bandwidth of the composite link that carries connections.

The solution shall support a policy-based preemption capability such that, in the event of such a "bandwidth shortage", the signaled or configured preemption and holding parameters can be applied to the following treatments to the connections:

- o For a connection that has RSVP-TE LSPs, signal the router that the LSP has been preempted. The solution shall support soft preemption (i.e., notify the preempted LSP source prior to preemption). [Soft Preemption]
- o For a connection that has LDP(s), where the routers connected via the composite link is aware of the LDP signaling involved to the preempted label stack depth, signal release of the label to the router
- o For a connection that has non-re-routable RSVP-TE LSPs or non-releasable LDP labels, signal the router or operator that the LSP or LDP label has been lost.

[4.2.](#) Exterior Functions

[4.2.1.](#) Functions common to all LSP flows

The solution MUST be able to interoperate with exiting IETF MPLS [[RFC3031](#)] control and data planes where appropriate.

[4.2.1.1.](#) Signaling Protocol Extensions

The solution SHALL support signaling a composite link between two routers (e.g., P, P/PE, or PE).

The solution SHALL support signaling a component link as part of a composite link.

The solution SHALL support signaling a composite link and automatically injecting it into the IGP [LSP Hierarchy bis] or connected two routers.

[4.2.1.2.](#) Routing Advertisement Extensions

The solution SHALL support IGP extensions to identify that the advertised routing adjacency as a composite link.

[4.2.1.3.](#) Multi-Layer Networking Aspects

The solution MUST be able to interoperate with existing IETF GMPLS/MPLS-TP control and data planes where appropriate, for example, when signaling a component link.

The solution SHALL support derivation of the advertised interface

parameters from signaled component link parameters from a lower layer (e.g., latency, capacity) based on operator policy.

The solution SHALL be able to accept the GMPLS/MPLS-TP control plane signaled component link. It SHALL be able to support the following component link parameters

- o Maximum (estimated or measured) acceptable latency
- o Actual(estimated or measured) assigned latency
- o Bandwidth
- o Delay Variation

It SHOULD support the following component link parameter

- o Loss rate

[4.2.2.](#) Functions specific to LSP flows with TE information

[4.2.2.1.](#) Signaling Protocol Extensions Requirements

The solution SHALL support per LSP signaling of at least the following additional parameters for an LSP

- o Maximum (estimated or measured) acceptable latency
- o Actual (estimated or measured) accumulated latency based upon the actual component link assigned by the composite link
- o Bandwidth of the highest and lowest speed

The solution SHOULD support per LSP signaling of at least the following additional parameters:

- o Delay Variation
- o Loss rate

As part of determining the path of an LSP, the client may query a Patch Computation Element (PCE) and use the response in determining the (complete or partial) source route sent in the TE signaling message.

Note that an LSP can be a component link or a client LSP. Since component links may involve GMPLS, separate signaling protocol extensions may be needed for particular switching capabilities.

4.2.2.2. Routing Advertisement Extensions

It SHALL be possible for the solution to represent multiple values, or a range of values, for the composite link interface parameters in order to communicate information about differences in the constituent component links in an exterior function route advertisement

Some of these capabilities are specified for link bundling [RFC 4201], but some extensions may be needed. Techniques such as those described in [[RFC 2676](#)] for encoding latency and using it in routing should be considered. LSA encoding techniques such as those described in [[RFC 3630](#)] should also be considered.

For example, a range of latencies, a range of capacities and/or other characteristics for the component links that comprise the composite links could be advertised. When a range of characteristics is advertised, these can be used in a constrained path computation, that is, certain composite links can be excluded since no component link meets the characteristic as part of the overall path.

4.2.3. Functions specific to LSP flows without TE information

A straightforward set of requirement would result if the same functions specified for RSVP-TE were also specified for LDP. However, the IETF MPLS Working Group's decision on MPLS signaling protocols [[RFC 3468](#)] to not pursue such an approach by not adding functions to CR-LDP means that a different approach must be taken.

As described in the Interior Function section, 4.1.3, a basic composite link capability when there is no TE information is to be able to measure the amount of LDP signaled traffic that is sent on an LSP. It is highly desirable to be able to signal and advertise certain aspects of these measurements, if they cannot be explicitly signaled via extensions to LDP.

4.2.3.1. Signaling Protocol Extensions

The solution SHOULD allow addition of an object to an LDP label mapping message that indicates how much measured capacity can be sent to a pair of nodes connected via a composite link. This signaling should be conveyed at least to nodes which are adjacent to the pair of nodes connected via a composite link.

Nodes SHOULD be able to use this information in conjunction with the IGP link information to decide which label mappings it will use for forwarding LDP signaled flows toward a pair of nodes connected via a

composite link.

4.2.3.2. Routing Advertisement Extensions

Signaling via LDP the available capacity on a composite link for flows without TE information is one method. A preferable method would be to extend the IGP to indicate the amount of capacity allocated by an Interior function to flows without TE information so that nodes in the network using LDP can make better informed decisions about which label mapping messages are used to form an LDP LSP.

[4.2.4.](#) Functions specific to LSP flows with and without TE information

As described in the Interior Function section, 4.1.4, the principal driver in this case is how to handle bandwidth shortage events when both RSVP-TE and LDP signaled LSPs are present in the composite link. The requirements relevant to the nodes connected by the composite link are described in [section 4.1.4](#), and this section describes additional desirable functions beyond these nodes. Since RSVP-TE signaling defines parameters and procedures for preemption, the focus is on extensions to better support LDP signaled LSPs.

[4.2.4.1.](#) Signaling Protocol Extensions

The solution SHOULD allow addition of an object to an LDP label mapping message that indicates that a node controlling the composite link wants to preempt the traffic offered by an LSP. This should have the effect of pruning label distribution for the LSP at the node pair connected via a composite link.

[4.2.4.2.](#) Routing Advertisement Extensions

The solution SHOULD allow addition of an object to the IGP that indicates that all LDP signaled traffic should avoid the advertised composite link.

[5.](#) Security Considerations

The solution is a local function on the router to support traffic engineering management over multiple parallel links. It does not introduce a security risk for control plane and data plane.

The solution could change the frequency of routing update messages and therefore could change routing convergence time. The solution MUST provide controls to dampen the frequency of such changes so as

to not destabilize routing protocols.

[6.](#) IANA Considerations

IANA actions to provide solutions are for further study.

So, et al.

Expires April 16, 2010

[Page 16]

7. References

[7.1.](#) Normative References

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

[RFC2215] S. Shenker, J. Wroclawski, "General Characterization Parameters for Integrated Service Network Elements." September 1997

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

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

[RFC4090] Pan, P., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", [RFC 4090](#), May 2005.

[RFC4124] Protocol Extensions for Support of Diffserv-aware MPLS Traffic Engineering F. Le Faucheur, Ed. June 2005

[RFC4201] Kompella, K., "Link Bundle in MPLS Traffic Engineering", [RFC 4201](#), March 2005.

[7.2.](#) Informative References

[Entropy Label] Kompella, K. and S. Amante, "The Use of Entropy Labels in MPLS Forwarding", November 2008, Work in Progress

[LSP Hierarchy] Shiimoto, K. and A. Farrel, "Procedures for Dynamically Signaled Hierarchical Label Switched Paths", November 2008, Work in Progress

[Soft Preemption] Meyer, M. and J. Vasseur, "MPLS Traffic Engineering Soft Preemption", February 2009, Work in Progress

[CLFRAMEWORK] Yong, L. Ed, "Framework for MPLS Over Composite Link," work in progress.

[RFC3468] L. Andersson, G. Swallow, "The Multiprotocol Label Switching (MPLS) Working Group decision on MPLS signaling protocols."

[PWBONDING] Stein, Y, "PW Bonding", Dec. 2008

So, et al.

Expires April 16, 2010

[Page 17]

[[RFC3630](#)] Katz, D., "Traffic Engineering (TE) Extension to OSPF Version 2", [RFC 3630](#), September 2003.

[RFC 2676] G. Apostolopoulos, S. Kama, D. Williams, R. Guerin, A. Orda, T. Przygienda, "QoS Routing Mechanisms and OSPF Extensions," August, 1999

[8](#). Acknowledgments

Authors would like to thank Adrian Farrel from Olddog for his extensive comments and suggestions, Ron Bonica from Juniper, Nabil Bitar from Verizon, Eric Gray from Ericsson, Lou Berger from LabN, and Kireeti Kompella from Juniper, for their reviews and great suggestions.

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

Copyright (c) 2009 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

So Ning
Verizon
2400 N. Glem Ave.,
Richerdson, TX 75082
Phone: +1 972-729-7905
Email: ning.so@verizonbusiness.com

Andrew Malis
Verizon
117 West St.
Waltham, MA 02451
Phone: +1 781-466-2362
Email: andrew.g.malis@verizon.com

Dave McDysan
Verizon
22001 Loudoun County PKWY
Ashburn, VA 20147
Email: dave.mcdysan@verizon.com

Lucy Yong
Huawei USA
1700 Alma Dr. Suite 500
Plano, TX 75075
Phone: +1 469-229-5387
Email: lucyyong@huawei.com

Frederic Jounay
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex,
FRANCE
Email: frederic.jounay@orange-ftgroup.com

Yuji Kamite
NTT Communications Corporation
Granpark Tower
3-4-1 Shibaura, Minato-ku
Tokyo 108-8118
Japan
Email: y.kamite@ntt.com