Routing Working Group N. So
Internet Draft A. Malis
Intended Status: Informational D. McDysan
Expires: Verizon
L. Yong
Huawei
F. Jounay
France Telecom
Y. Kamite
NTT
July 9, 2009

Framework and Requirements for MPLS Over Composite Link
draft-so-yong-mpls-ctg-framework-requirement-02

Status of this Memo

Copyright Notice

Abstract

This document states a traffic distribution problem in today's IP/MPLS network when multiple links are configured between two routers. The document presents motivation, a framework and requirements. 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 document primarily focuses on MPLS traffic controlled through control plane protocols, the advertisement of composite link parameter in routing protocols, and the use of composite links in the RSVP-TE and LDP signaling protocols. Interactions with the data and management plane are also addressed. Applicability can be between a single pair of MPLS-capable nodes, a sequence of MPLS-capable nodes, or a multi-layer network connecting MPLS-capable nodes.

Table of Contents

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.

This document also describes a framework for selecting members of a Composite Link, operating the Composite Link in signaling and routing, and for distributing through local decisions data flows across the component members of a Composite Link to achieve maximal data throughput and enable link-level protection schemes.

Applicability of the work within this document is focused on MPLS traffic as controlled through control plane protocols. Thus, this document describes 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 framework and requirements of this document.

The transport functions for TE and non-TE traffic delivery over a Composite Link are termed a Composite Transport Group (CTG). In other words, the objective of CTG is to solve the traffic sharing problem at a composite link level by mapping labeled traffic flows to component links:

1. using TE information from the control plane attached to the virtual interface when available, or

2. using traffic measurements when it is not.

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

CTG: Composite Transport Group

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 or Composite Transport Group (CTG): 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.)

CTG 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.

CTG 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.

CTG 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.

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

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.

- 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.

- 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).

- 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 CTG 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.

3.2. Summary of Problems Requiring Solution

The following bullets highlight aspects of CTG-related solution for which detailed requirements are stated in Section 5.

- 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.

- Extend a link bundling type function to scenarios with groups of links having different characteristics (e.g., bandwidth, latency).

- When an end-to-end LSP signaled by RSVP-TE uses a composite link, the CTG must select a component link that meets the end-to-end requirements for the LSP. To perform this function, the CTG must be made aware of the required, desired, and acceptable link characteristics (e.g., latency, optimization frequency) for each CTG hop in the path.

- 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).

- In the case where multiple routing instances could share a composite link, a means to reduce or manage the potential inefficiency is highly desirable. A local implementation by the same router type at each end of a CTG could address this issue. However, in the case of different routers at each end of a CTG there is a need to specify the operational configuration commands and measurements to ensure interoperability. Alternatively, the case of multiple routing instances sharing a CTG could be viewed as an instance of multi-layer routing. In this case, some lower-layer instance of routing associated with the CTG can be viewed as a server. This CTG server controls the composite link and arbitrates between the signaled requests and measured load offered by the higher level, client instances of routing (i.e., users of the CTG). The CTG server assigns resources on component links to these client level routing instances and communicates this via routing messages into each of the client instances, which then communicate this to their peers in the domain of each routing instance. This server level function is a way to meet operational requirements where flows from one routing instance need to preempt flows from another routing instance, as detailed in the requirements in section 6.1.3.

4. Framework

4.1. Single Routing Instance

4.1.1. Summary Block Diagram View

The CTG framework for a single routing instance is illustrated in Figure 1, where a composite link is configured between routers R1 and R2. In this example, the composite link has three component links. A composite link is defined in ITU-T [ITU-T G.800] as a single link that bundle comprises multiple parallel component links between the two routers. Each component link in a composite link is supported by a separate server layer trail. A component link can be implemented by different transport technologies such as wavelength, SONET/SDH, OTN, Ethernet PHY, Ethernet VLAN, or can be a logical link [LSP Hierarchy] for example, MPLS, or MPLS-TP. Even if the transport technology implementing the component links is identical, the characteristics (e.g., bandwidth, latency) of the component links may differ.

An important framework concept is that of a CTG connection shown in Figure 1. Instead of simply mapping the incoming traffic flows directly to the component links, aggregating multiple flows into a connection makes the measurement of actual bandwidth usage more scalable and manageable. Then the CTG can place connections in a 1:1 manner onto the component links. Although the mapping of flows to connections and then to a component link is a local implementation matter, the management plane configuration and measurement of this mapping is an important external operational interface necessary for interoperability. Note that a special case of this model is where a single flow is mapped to a single connection.

```
            Management Plane
        Configuration and Measurement <----------------+
                    ^                                  |
                    |                                  |
                    |                                  |
                    v                                  v
                +---------+                           +---------+
        Control |      R1 |                           |      R2 |   Control
        Plane ====>       |                           |       ====> Plane
                |     +---+     Component Link 1      +---+     |
                |     |   |===========================|   |     |
                |     |   |~~~~~~ CTG Connections ~~~~|   |     |
              ~~|~~>~~|   |===========================|   |~~>~~|~~
              ~~|~~>~~| C |    Component Link 2       | C |~~>~~|~~
      Traffic ~~|~~>~~|   |===========================|   |~~>~~|~~ Traffic
      Flows   ~~|~~>~~| T |~~~~~~ CTG Connections ~~~~| T |~~>~~|~~  Flows
              ~~|~~>~~|   |===========================|   |~~>~~|~~
              ~~|~~>~~| G |     Component Link 3      | G |~~>~~|~~
              ~~|~~>~~|   |===========================|   |~~>~~|~~
                |     |   |~~~~~~ CTG connections ~~~~|   |     |
                |     |   |===========================|   |     |
                |     +---+                           +---+     |
                +---------+                           +---------+
                        ! !                           ! !
                        ! !<---- Component Links ---->! !
                        !<------ Composite Link ------->!
```

Figure 1: Composite Transport Group Architecture Model

CTG functions can be grouped into two major categories, as described in the following subsections.

4.1.2. CTG Interior Functions

CTG Interior Functions: implemented within the interior of MPLS routers connected via a composite link. This includes the local data plane functions of determining the component link on which a traffic flow is placed. Management configuration for some aspects of these interior functions is important to achieve operational consistency and this is the focus of requirements in this document for interior functions.

4.1.3. CTG Exterior Functions

CTG Exterior Functions have aspects that are applicable exterior to the CTG connected MPLS routers. In other words, functions that are used by other routers, such as routing advertisements and signaling messages related to specific characteristics of a composite link.

4.1.4. Multi-Layer Network Context

The model of Figure 1 applies to at least the scenarios illustrated in

Figure 2. The component links may be physical or logical, and the composite link may be made up of a mixture of physical and logical links supported by different technologies. Figure 2 and the following description provide a contextual framework for the multi-layer

networking related problem described in section 3.2. In the first scenario, a set of physical links connect adjacent (P) routers (R1/R2).

In the second scenario, a set of logical links connect adjacent (P or PE) routers over other equipment (i.e., R3/R4) that may implement RSVP-TE signaled MPLS tunnels which may be in the same IGP as R1/R2 or in a different IGP. . When R3 and R4 are not part of R1/R2's IGP (e.g., they may implement MPLS-TP) R3/R4 can have a signaling but not a routing interface with R1/R2. In other words, R3/R4 offers connectivity to R1/R2 in an overlay model. Another case is where R3/R4 provide a TE-LSP segment of TE-LSP from R1 and R2.

```
 +----+---+                 1. Physical Link            +---+----+
 |    |   |----------------------------------------------|   |    |
 |    |   |                                              |   |    |
 |    |   |     +------+                      +------+   | C |    |
 |    | C |     | MPLS |    2. Logical Link   | MPLS |   |   |    |
 |    |   |.... |......|......................|......|....|   |    |
 |    |   |-----|  R3  |----------------------|  R4  |----|   |    |
 |    | T |     +------+                      +------+   | T |    |
 |    |   |                                              |   |    |
 |    |   |                                              |   |    |
 |    | G |     +------+                      +------+   | G |    |
 |    |   |     |GMPLS | 3. Lower Layer Link  |GMPLS |   |   |    |
 |    |   |. ...|......|......................|......|....|   |    |
 |    |   |-----|  R5  |----------------------|  R6  |----|   |    |
 |    |   |     +------+                      +------+   |   |    |
 | R1 |   |                                              |   | R2 |
 +----+---+                                              +---+----+
           |<---------- Composite Link ----------------->|
```

Figure 2: Illustration of Component Link Types

In the third scenario, GMPLS lower layer LSPs (e.g., Fiber, Wavelength, TDM) as determined by a lower layer network in a multi-layer network deployment as illustrated by R5/R6. In this case, R5 and R6 would usually not be part of the same IGP as R1/R2 and may have a static interface, or may have a signaling but not a routing association with R1 and R2. Note that in scenarios 2 and 3 when the intermediate routers are not part of the same IGP as R1/R2 (i.e., can be viewed as operating at a lower layer) that the characteristics of these links (e.g., latency) may change dynamically, and there is an operational desire to handle this type of situation in a more automated fashion than is currently possible with existing protocols. Note that this problem currently occurs with a single lower-layer link in existing networks and it would be desirable for the solution to handle the case of a single lower-layer component link as well. Note that the interfaces at R1 and R2 are associated with these different component links can be configured with IP addresses or use unnumbered links as an interior, local function since the individual component links are not advertised as the CTG virtual interface.

4.2. Multiple Routing Instances

In the case where the routers connected via a CTG support multiple routing instances there is additional context as described in this section. In general, each routing instance can have its own instances of control plane, IGP, and/or routing/signaling protocols. In general, they need not be aware of the existence of the other routing instances. However, it is operationally desirable for efficiency reasons for these routing instances to share the resources of a composite link and have the capability for a higher level of control logic to allocate resources amongst the instances based upon configured policy and the current state of at least the local composite link, but potentially that of other composite links in the network. Figure 3 shows the model where a composite link appears as a routable virtual interface to each routing instance.

```
    +-----+---+       Component Link1                  +---+-----+
    |     |   |----------------------------------------|   |     |
    |RIA.1|   |                                        |   |RIA.2|
    |     | C |        Virtual Interface               | C |     |
    |IGPA====================================================IGPA|
    |_____| T |       Component Link2                  | T |_____|
    |     |   |----------------------------------------|   |     |
    |RIB.1| G |                                        | G |RIB.2|
    |     |   |       Component Link3                  |   |     |
    |     |   |----------------------------------------|   |     |
    |IGPB====================================================IGPB|
    +-----+---+        Virtual Interface               +---+-----+
          |                                                |
          |<------------- Composite Link ----------------->|
            Figure 3: Routing Instances Sharing Composite Link
```

In Figure 3, the router on the left side is configured with two routing instances (RI) RIA.1 and RIB.1. Another router on the right side is configured with two routing instances RIA.2 and RIB.2. Routing instance A belongs to IGPA network and routing instance B belongs to IGPB network. In this example the composite link contains three component links. IGPA and IGPB can be TE and/or non-TE enabled. In this case, there are additional CTG related functions related to the dynamic allocation of resources in the component links to each of the multiple routing instances. Furthermore, there are operational scenarios where in response to certain failure scenarios and/or load conditions that the multi-routing instance CTG function may preempt certain LSPs and/or cause changes in the routing information communicated by the IGPs as detailed in the section on multi-instance CTG exterior function requirements.

The multiple routing instance case of CTG appears to have a number of requirements and context in common with the single routing instance of

CTG, and hence it is retained within the same document in this version. The structure of this framework section, as well as the following requirements section, is to place the multiple routing instance CTG requirements at the end and to only describe aspects unique to the multiple routing instance case.

The larger view of CTG as a higher level instance in the context of multiple lower level routing instances may be sufficiently different and broad enough in scope to justify elaboration in a separate document. However, an objective should be to use the framework and as many common requirements from the single routing instance CTG framework and requirements as possible.

5. CTG Requirements for a Single Routing Instance

5.1. Management and Measurement of CTG Interior Functions

5.1.1. Configuration as a Routable Virtual Interface

The operator SHALL be able to configure a "virtual interface" corresponding to a composite link and component link characteristics as a TE link or an IP link in IP/MPLS network.

The solution SHALL allow configuration of virtual interface parameters for a TE 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 allow configuration of virtual interface parameters for an IP link used for MPLS (e.g., administrative cost or weight).

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

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

The "virtual interface" SHALL appear as a fully-featured routing adjacency in each routing instance, 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.

CTG 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.

5.1.2. Traffic Flow and CTG Mapping

CTG SHALL support operator assignment of traffic flows to specific connections.

CTG SHALL support operator assignment of connections to specific component links.

CTG 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).

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

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

5.1.2.1. Using Control Plane TE Information

The following requirements apply to the case of RSVP-TE signaled LSPs.

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

CTG 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.

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

o Local Bandwidth Oversubscription factor

5.1.2.2. When no TE Information is Available (i.e., LDP)

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

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

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 flows).

5.1.2.3. Handling Bandwidth Shortage Events

The following requirements apply to a virtual interface that supports the traffic flows both with and without TE information, in response to a bandwidth shortage event. A "bandwidth shortage" can arise in CTG 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 the CTG connections.

CTG 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. CTG 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 CTG 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.

5.1.3. Management of Other Operational Aspects

5.1.3.1. Resilience

Component links in a composite link may fail independently. The failure of a component link may impact some CTG connections. The impacted CTG connections shall be transferred to other active component links using the same rules as for the original assignment of CTG connections to component links.

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.

5.1.3.2. 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 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

o A manually configured setting at composite link ingress

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

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

5.1.3.3. OAM Messaging Support

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.

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)
- o Control-plane of the virtual interface level (i.e., routing/signaling on it)
- o o A CTG that believes that the underlying server layer might not efficiently report failures, can run Bidirectional Forwarding Detection (BFD) over a component link.

CTG shall support configuration of timers so that lower layer methods have time to detect/restore faults before a CTG function would be invoked.

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

5.2. CTG Exterior Functions

5.2.1. 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 or a private link for connected two routers.

The solution SHALL support signaling of at least the following additional parameters for component link:

- o Minimum and Maximum (estimated or measured) latency
- o Bandwidth of the highest and lowest speed

The solution SHOULD support signaling of at least the following additional parameters for component link:

- o Delay Variation
- o Loss rate

5.2.2. Routing Advertisement Extensions

It shall be possible 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. For example, a range of latencies for the component links that comprise the composite links could be advertised.

Multi-Layer Networking Aspects

The solution SHALL support derivation of the advertised interface parameters from signaled component link parameters from a lower layer (e.g., latency) based on operator policy.

6. CTG Requirements for Multiple Routing Instances

This section covers requirements conditioned on the case where the solution supports multiple routing instances. Unless otherwise stated, all requirements for a single routing instance from section 5 apply individually to each of the multiple routing instances.

6.1. Management and Measurement of CTG Interior Functions

6.1.1. Appearance as Multiple Routable Virtual Interfaces

CTG SHALL support multiple routing instances that see a single separate "virtual interface" to a shared composite link composed of parallel physical/logical component links between a pair of routers.

6.1.2. Control of Resource Allocation

The operator SHALL be able to statically assign resources (e.g., component link, or bandwidth to a sub/logical interface) to each routing instance virtual interface.

6.1.3. Configuration of Prioritization and Preemption

The solution SHALL support a policy based local to the CTG preemption capability across all routing instances and a set of requirements similar to those listed in section 5.1.2.3. Note that this requirement applies across the multiple routing instances.

6.2. CTG Exterior Functions

6.2.1. CTG Operation as a Higher-Level Routing Instance

The following requirements apply to the case where CTG exterior functions supporting multiple routing instances communicate with each other.

CTG exterior functions shall be able to advertise parameters such as reserved capacity, measured capacity usage, and available resources for the CTGs of which they perform CTG interior functions.

CTG exterior functions shall be able to signal and respond to requests for a change in allocation of the CTG interior function resources.

7. 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.

8. IANA Considerations

IANA actions to provide solutions are for further study.

9. References

9.1. 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

[RFC3477] Kompella, K., "Signalling Unnumbered Links in Resource ReSerVation Protocol - Traffic Engineering (RSVP-TE)", RFC 3477, January 2003.

[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.

9.2. Informative References

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

[LSP Hierarchy] Shiomoto, 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

10. 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.

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