Network Working Group
Internet Draft
Proposed Status: Informational
Expires: April 2007

Tomonori Takeda
NTT
Yuichi Ikejiri
NTT Communications
Adrian Farrel
Old Dog Consulting
Jean-Philippe Vasseur
Cisco Systems, Inc.

October 2006

Analysis of Inter-domain Label Switched Path (LSP) Recovery
draft-takeda-ccamp-inter-domain-recovery-analysis-01.txt

Status of this Memo

Abstract

This document analyzes various schemes to realize Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) Label Switched Path (LSP) recovery in multi-domain networks based on the existing framework for multi-domain LSPs.

The main focus for this document is on establishing end-to-end diverse Traffic Engineering (TE) LSPs in multi-domain networks. It presents various diverse LSP setup schemes based on existing functional elements.

Table of Contents

1. Terminology..2
2. Introduction...3
2.1 Domain..3
2.2 Document Scope..4
2.3 Note on Other Recovery Techniques.................................5
2.4 Signaling Options...5
3. Diversity in Multi-domain Networks.................................5
3.1 Multi-domain Network Topology.....................................6
3.2 Note on Domain Diversity..7
4. Factors to Consider..7
4.1 Scalability versus Optimality.....................................7
4.2 Key Concepts..8
5. Diverse LSP Setup Schemes without Confidentiality.................10
5.1 Management Configuration...10
5.2 Head-end Path Computation (with multi-domain visibility).....10
5.3 Per-domain Path Computation......................................10
5.3.1 Sequential Path Computation....................................11
5.3.2 Simultaneous Path Computation..................................11
5.4 Inter-domain Collaborative Path Computation.....................12
5.4.1 Sequential Path Computation....................................12
5.4.2 Simultaneous Path Computation..................................13
6. Diverse LSP Setup Schemes with Confidentiality....................13
6.1 Management Configuration...14
6.2 Head-end Path Computation (with multi-domain visibility).....14
6.3 Per-Domain Path Computation......................................15
6.3.1 Sequential Path Computation....................................15
6.3.2 Simultaneous Path Computation..................................15
6.4 Inter-domain Collaborate Path Computation.......................16
6.4.1 Sequential Path Computation....................................16
6.4.2 Simultaneous Path Computation..................................16
7. Network Topology Specific Considerations..........................17
8. Addressing Considerations...17
9. Note on SRLG Diversity..17
10. Manageability Considerations.....................................17
11. Security Considerations..18
12. References...18
12.1 Normative References..18
12.2 Informative References..18
13. Acknowledgments..20
14. Author's Addresses...20
15. Intellectual Property Consideration..............................20
16. Full Copyright Statement...21

1. Terminology

The reader is assumed to be familiar with the terminology in [inter-domain-fw] that provides a framework for inter-domain Label Switched Path (LSP) setup, and [RFC4427] that provides terminology for LSP recovery.

The following are several key terminologies used within this document.

- Domain: See [inter-domain-fw]. A domain is considered to be any collection of network elements within a common sphere of address management or path computational responsibility. Note that nested domains continue to be out of scope.

- Working LSP: See [RFC4427]. The working LSP transports normal user traffic. Note that the term LSP and TE LSP will be used interchangeably.

- Recovery LSP: See [RFC4427]. The recovery LSP transports normal user traffic when the working LSP fails. The recovery LSP may transport user traffic even when the working LSP is transporting normal user traffic (e.g., 1+1 protection). In such a scenario, the recovery LSP is sometimes referred to as a protecting LSP.

- Diversity: See [inter-domain-fw]. Diversity means the relationship of multiple LSPs, where those LSPs do not share some specific type of resource (e.g., link, node, or shared risk link group (SRLG)). Diverse LSPs may be used for various purposes, such as load-balancing and recovery. In this document, the recovery aspect of diversity, specifically the end-to-end diversity of two traffic engineering (TE) LSPs, is the focus. Those two diverse LSPs are referred to as the working LSP and recovery LSP hereafter. Sometimes, diversity is referred to as disjointness.

- Confidentiality: See [RFC4216]. The term confidentiality applies to the protection of information about the topology and resources present within one domain from visibility by people or applications outside that domain.

2. Introduction

2.1 Domain

As defined in [inter-domain-fw], a domain is considered to be any collection of network elements within a common sphere of address management or path computational responsibility. Examples of such domains include IGP areas and Autonomous Systems. A network accessed over the Generalized Multiprotocol Label Switching (GMPLS) User-to-Network Interface (UNI) [RFC4208] and a Layer One Virtual Private

Network (L1VPNs) [L1VPN-FW] are special cases of multi-domain networks.

Example objectives of domain usage include administrative separation, scalability, and forming protection domains.

As described in [inter-domain-fw], there could be TE parameters (such as color and priority) whose meanings are specific to each domain. In such a scenarios, mapping functions could be performed based on policy agreements between domain administrators.

2.2 Document Scope

This document analyzes various schemes to realize Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) LSP recovery in multi-domain networks based on the existing framework for multi-domain LSP setup [inter-domain-fw].

There are various recovery techniques for LSPs. For TE LSPs, major techniques are end-to-end recovery [e2e-recovery], local protection such as Fast Reroute (FRR) [RFC4090] (in packet switching environments), and segment recovery [seg-recovery] (in GMPLS).

In this version of the document the main focus is the analysis of diverse TE LSP (hereafter LSP) setup schemes, which can advantageously used in the context of end-to-end recovery. This document presents various diverse LSP setup schemes by combining various functional elements. Analysis of other recovery techniques could be added in a later revision of this document if necessary. Furthermore, details of maintenance of diverse LSPs, such as re-optimization and OAM, are for further study.

Note that the comparative evaluation of these various schemes is out of scope for this document, and should be described in separate applicability documents.

[RFC4105] and [RFC4216] describe requirements for diverse LSPs. There could be various types of diversity, and this document focuses on node/link/SRLG diversity. Note that domain diversity (that is, the selection of paths that have only the ingress and egress domains in common) is discussed in section 3.2.

Based on the service grade, both the working LSP and the recovery LSP may be established at the time of service establishment (e.g., service requiring high resiliency). Alternatively, the recovery LSP may be added later due to a change in the grade of the service.

Also, the recovery LSP may be used for 1+1 protection, 1:1 protection, or shared mesh restoration. However, ways to compute

diverse paths, and the signaling of these TE LSPs are common across all uses, and these topics are the main scope of this document.

Section 5 of [inter-domain-fw] describes some analysis of diverse LSPs in multi-domain networks, and this document provides more detailed analysis based on that content.

Note that diverse LSPs may be used for various purposes, in addition to recovery. An example is for load-balancing, so as to limit the traffic disruption to a portion of the traffic flow in case of a single network element failure. This document does not preclude use of diverse LSP setup schemes for those purposes.

2.3 Note on Other Recovery Techniques

Fast Reroute and segment recovery in multi-domain networks are described in section 5.4 of [inter-domain-fw], and a more detailed analysis is provided in section 5 of [inter-domain-rsvp]. Additional analysis may be added in a future revision of this document if necessary.

Also, the recovery type of an LSP or service may change at a domain boundary. That is, the recovery type would remain the same within one domain, but might be different in the next domain. This may be due to the capabilities of each domain, administrative policies, or to topology limitations. An example is where protection at the domain boundary is provided by link protection on the inter-domain link(s), but where protection within each domain is achieved through segment recovery. This mixture of protection techniques is for further study.

2.4 Signaling Options

There are three signaling options for establishing inter-domain TE LSPs: nesting, contiguous LSPs, and stitching [inter-domain-fw]. The description in this document of diverse LSP setup is agnostic in relation to the signaling option used, unless otherwise specified.

Note that signaling option specific considerations for Fast Reroute and segment recovery are described in section 5 of [inter-domain-rsvp]. Also note that if the recovery type changes at the domain boundary, the nesting and stitching options may be more suitable. Details are for further study.

3. Diversity in Multi-domain Networks

As described in section 2.2, analysis of diverse LSP setup schemes in multi-domain networks is the main focus of this document. This section describes some assumptions in this problem space made in this document.

3.1 Multi-domain Network Topology

Figures 1 and 2 show example multi-domain network topologies. In Figure 1, domains are connected in a linear topology. There are multiple paths between nodes A and L, but all of them cross domain#1-domain#2-domain#3 in that order.

```
     +--Domain#1--+   +--Domain#2--+   +--Domain#3--+
     |            |   |            |   |            |
     |     B------+---+---D-----E--+---+------J     |
     |    /       |   |    \   /   |   |       \    |
     |   /        |   |     \ /    |   |        \   |
     |  A         |   |      H     |   |         L  |
     |   \        |   |     / \    |   |        /   |
     |    \       |   |    /   \   |   |       /    |
     |     C------+---+---F-----G--+---+------K     |
     |            |   |            |   |            |
     +------------+   +------------+   +------------+
```

Figure 1: Linear Connectivity

```
                     +-----Domain#2-----+
                     |                  |
                     | E--------------F |
                     | |              | |
                     +-+--------------+-+
                       |              |
            +-Domain#1-+-+  +---------+--+
            |          | |  |         |  |
            |  A-------B-+--+-C-------D  |
            |  |         |  | |          |
            +--+---------+  +-+-Domain#4-+
               |              |
             +-+--------------+-+
             | |              | |
             | G--------------H |
             |                  |
             +-----Domain#3-----+
```

Figure 2: Mesh Connectivity

In Figure 2, domains are connected in a mesh topology. There are multiple paths between nodes A and D, and these paths do not necessarily follow the same set of domains.

Indeed, if inter-domain connectivity forms a mesh, domain level routing is required (even for the unprotected case). This is tightly coupled with the next hop domain resolution/discovery mechanisms.

In this version of the document, we assume that domain level routing is given via configuration or policy, and this is not part of path computation process described later in this document. Details on more advanced domain level routing are for further study.

In addition, domain level routing may perform "domain re-entry", where a path enters a domain after the path exits that domain (e.g., domain#X-domain#Y-domain#X). This requires specific considerations when confidentiality is required (described in section 4.2), and is for further study.

Furthermore, the working LSP and the recovery LSP may or may not be routed along the same set of domains in the same order. In this version of the document, we assume that the working LSP and recovery LSP follow the same set of domains in the same order (via configuration or policy). Details on other scenarios are for further study.

3.2 Note on Domain Diversity

As described in section 2.2, domain diversity means the selection of paths that have only the ingress and egress domains in common. This may provide enhanced resilience against failures, and is a way to ensure path diversity for most of the path of the LSP.

In section 3.1 we assumed that the working LSP and the recovery LSP follow the same set of domains in the same order. Under this assumption, domain diversity cannot be achieved. However, by relaxing this assumption, domain diversity could be achieved, e.g., by either of the following schemes.

- Consider domain diversity as a special case of SRLG diversity (i.e., assign an SRLG ID to each domain)
- Configure domain level routing to provide domain diverse paths (e.g., by means of AS_PATH in BGP)

Details are for further study, should it been considered as a requirement.

4. Factors to Consider

4.1 Scalability versus Optimality

As described in [inter-domain-fw], scalability and optimality are two conflicting objectives. Note that the meaning of optimality differs depending on each network operation. Some examples of optimality in the context of diverse LSPs are:

- Minimizing the sum of their cost while maintaining diversity
- Restricting the difference of their cost (so as to minimize delay difference after switch-over) while maintaining diversity

By restricting TE information distribution to only within each domain (and not across domain boundaries) as required by RFC4105 and RFC4216, it may not be possible to compute an optimal path. As such, it may not be possible to compute diverse paths, even if they exist. However, if we assume domain level routing is given (as discussed in section 3), it is possible to compute diverse paths in some schemes if such paths exist. This is discussed in section 5.

4.2 Key Concepts

Three key concepts to classify various diverse LSP setup schemes are presented below.

o With or without confidentiality

  - Without confidentiality

    Under this network configuration, it is possible to specify (by means of the Explicit Route Object - ERO comprising a list of strict hops) or obtain (by means of the recorded Route Object - RRO) a route across other domains.

    Examples of this configuration are multi-area networks, and some forms of multi-AS networks (especially within the same provider).

  - With confidentiality

    Under this network configuration, it is not possible to specify or obtain a route (by means of ERO/RRO) across other domains. Paths may be specified/obtained in the form of ERO/RRO containing loose hops. Therefore, it is not possible to specify or obtain a full route at the head-end of a multi-domain LSP.

    Examples of this configuration are some forms of multi-AS networks (especially inter-provider networks), GMPLS-UNI networks, and L1VPNs.

o Per domain path computation or inter-domain collaborative path computation

  - Per domain path computation

    In this scheme, a path is computed domain by domain as LSP signaling progresses through the network. This scheme requires ERO expansion in each domain. The case for unprotected LSPs under

this scheme is covered in [inter-domain-pd].

- Inter-domain collaborative path computation

  In this scheme, path computation is typically done before signaling. This scheme typically uses communication between cooperating path computation elements (PCEs) [PCE-arch]. The case for unprotected LSP under this scheme is covered in [brpc].

Note that these are path computation techniques. It is also possible to obtain a path via management configuration, or head-end path computation (with multi-domain visibility). This is also discussed in sections 5 and 6.

Note also that it is possible to combine multiple path computation techniques (including using a different technique for the working and recovery LSPs), but this is for further study and is likely to require sequential path computation (see below).

o Sequential path computation or simultaneous path computation

- Sequential path computation

  The path of the working LSP is computed (without considering the recovery LSP), and then the path of the recovery LSP is computed. Typically, this scheme is applicable when the recovery LSP is added later as change of the service grade. But this scheme can also be applicable when both of the working and recovery LSPs are established from the start. In this scheme, diverse LSPs may not be correctly computed (without some form of re-optimization or recomputation technique) in "trap" topologies. Furthermore, such sequential path computation approach may prevent from finding an "optimal" solution with regards to a specific objective function.

- Simultaneous path computation

  The path of the working LSP and the path of the recovery LSP are computed simultaneously. Typically, this scheme is applicable when both the working LSP and the recovery LSP are established together. In this scheme, it is possible to avoid trap topologies. Furthermore it may allow for achieving more optimal results.

Note that LSP setup with or without confidentiality is given as a per-domain configuration, while the choices of per-domain path computation or inter-domain collaborative path computation, and sequential path computation or simultaneous path computation may be a matter of choice for each individual pair of working/recovery LSPs.

The analysis of various diverse LSP setup schemes is described in sections 5 and 6, based on above criteria.

Some other considerations, such as network topology specific considerations, addressing considerations, and SRLG diversity are described in sections 7, 8 and 9.

5. Diverse LSP Setup Schemes without Confidentiality

In the following, various schemes for diverse LSP setup are presented based on different path computation techniques assuming that there is no requirement for confidentiality between domains. Section 6 makes a similar examination of schemes where inter-domain confidentiality is required.

5.1 Management Configuration

Section 3.1 of [inter-domain-fw] describes this path computation technique. The full explicit paths for the working and recovery LSPs are specified by a management application at the head-end, and no further computation or signaling specific considerations are needed.

5.2 Head-end Path Computation (with multi-domain visibility)

Section 3.2.1 of [inter-domain-fw] describes this path computation technique. The full explicit paths for the working and recovery LSPs are computed at the head-end either by the head-end itself or by a PCE. In either case the computing entity has full TE visibility across multiple domains and no further computation or signaling specific considerations are needed.

5.3 Per-domain Path Computation

Sections 3.2.2, 3.2.3 and 3.3 of [inter-domain-fw] describe this path computation technique. More detailed procedures are described in [inter-domain-pd].

In this scheme, the explicit paths of the working and recovery LSPs are specified as the complete strict path in the source domain followed by one of the following:

- The complete list of boundary LSRs (or domain identifiers, e.g., AS numbers) along the path.

- The boundary LSR for the source domain and the LSP destination.

Thus, ERO expansion is needed at domain boundaries. Path computation is performed by, or by a PCE on behalf of, each domain boundary LSR.

For establishing diverse LSPs using per-domain computation, there are two specific schemes, which are described in sections 5.3.1 and 5.3.2 respectively.

5.3.1 Sequential Path Computation

The Exclude Route Object (XRO) [XRO] can be used. Details are as follows.

Assume that the working LSP is established as described in [inter-domain-pd]. Also, assume that the route of the working LSP (full route) is available at the head-end through the RRO.

o Path computation aspect

  When performing path computation (ERO expansion) for the recovery LSP as it crosses each domain boundary a path is selected that avoids the nodes/links/SRLGs used by the working path so that a diverse path is obtained.

o Signaling aspect

  In order that the computation noted above can be performed, each computation point must be aware of the path of the working LSP. This information can be supplied in the XRO included in the Path message for recovery LSP. The XRO lists nodes, links and SRLGs that must be avoided by the LSP being signaled, and its contents are copied from the RRO of the working LSP.

This scheme cannot guarantee to establish diverse LSPs (even if they could exist) because the first LSP is established without consideration of the need for a diverse recovery LSP. Crankback [crankback] may be used in combination with this scheme in order to improve the possibility of successful diverse LSP setup. Furthermore, re-optimization of the working LSP and the recovery LSP may be used to achieve fully diverse paths.

Note that even if a solution is found, the degree of optimality of the solution (set of diverse TE LSP) might not be maximized.

5.3.2 Simultaneous Path Computation

o Path computation aspect

  When signaling the working LSP, the path of not only the working LSP, but also the recovery LSP is computed. For example, in Figure 1, when node D receives a Path message for the working LSP between nodes A and L, node D expands the ERO to reach domain#3. At the same time, node D computes a path for the recovery LSP across

the same domain from node F to reach domain#3. The entry boundary node for the recovery LSP (node F) needs to be specified in the Path message for the working LSP. In this example the path for the working LSP may be computed by node D as D-E-domain#3, and the path for recovery LSP as F-G-domain#3.

o Signaling aspect

There must be a mechanism to force the recovery LSP to follow the route computed above. One way to realize this is to have a specific object (with the same format as the ERO) to collect the route of the recovery LSP in the Path message for the working LSP and to return is to the head-end in the Resv message. When signaling the recovery LSP, the content of the ERO is copied from this object.

This scheme also cannot guarantee to establish diverse LSPs (even if they could exist) if there are more than two domain boundary nodes out of each domain. Crankback [crankback] may also be used in combination with this scheme to improve the chances of success.

Note that even if a solution is found, the degree of optimality of the solution (set of diverse TE LSP) might not be maximized.

5.4 Inter-domain Collaborative Path Computation

Section 3.4 of [inter-domain-fw] describes this approach. [brpc] provides some more detail. Path computation is performed for each of the working and recovery LSPs by the use of inter-PCE communication before each LSP is signaled.

There are two specific schemes for establishing diverse LSPs using this scheme, which are described in sections 5.4.1 and 5.4.2.

5.4.1 Sequential Path Computation

Route exclusion using the XRO [XRO] can be requested in the PCE communication protocol (PCEP) [PCEP] and this can be used to compute the path of a recovery LSP after the path of the working LSP has been determined. Details are as follows.

The working LSP is computed and may be immediately established as described in [brpc]. Assume that the path of the working LSP (full route) is available from the RRO.

o Path computation aspect

When requesting path computation for the recovery LSP, an XRO is included in the PCEP path computation request message (see [PCEP]). The content of the XRO is copied from the RRO of the working LSP.

The computation request specifies that the full route must be returned. Per-domain PCEs may need to be invoked by the first PCE that is consulted in order to collaboratively determine the correct path for the recovery LSP (just as described in [PCE-arch] and [inter-domain-fw] for the computation of a single inter-domain LSP by collaborating PCEs), and these PCEs exchange the XRO information to ensure that the computed path is diverse from the working LSP.

o Signaling aspect

The recovery LSP is signaled by including an ERO whose content is copied from the result returned by the PCE.

This scheme cannot guarantee to establish diverse LSPs (even if they exist) because the working LSP may be blocking. In such a scenario, re-optimization of the working and recovery LSPs may be used to achieve fully diverse paths.

Note that PCEP [PCEP] does not currently include support for the XRO, but that this is planned to be added in a future version.

5.4.2 Simultaneous Path Computation

o Path computation aspect

The PCEP SVEC Object allows diverse path computation to be requested. It would be possible to extend [brpc] to compute diverse paths. Details are for further study.

o Signaling aspect

In this scheme the PCE returns the full routes for the working and recovery LSPs and they are signaled accordingly.

This scheme can guarantee to establish diverse LSPs (if they exist), assuming domain level routing is given as described in section 3.

Furthermore, the computed set of TE LSPs may be optimal with respect to some objective functions.

6. Diverse LSP Setup Schemes with Confidentiality

In the context of this section, the term confidentiality applies to the protection of information about the topology and resources present within one domain from visibility by people or applications outside that domain. This includes, but is not limited to, recording of LSP routes, in addition to advertisements of TE information. The confidentiality does not apply to the protection of user traffic.

Diverse LSP setup schemes with confidentiality are similar to ones without confidentiality. However, several additional mechanisms are needed to preserve confidentiality. Examples of such mechanisms are:

- Path key: Provide each per-domain segment of the path in advance to the domain boundary nodes or to the PCE that computed the path for a limited period of time (temporary caching) and identify it in the signaled ERO using a path key. When path computation is done by PCE, the identify of the PCE containing state for the path may be required as well (e.g., PCE I-D). The path key is provided by the PCE to the head-end for inclusion in the ERO [conf-segment].

- LSP segment: Pre-establish each per-domain segments of the path using hierarchical LSPs [RFC4206] or LSP stitching segments [LSP-stitch] and signal the end-to-end LSP to use these per-domain LSPs. This scheme may need additional identifiers (such as LSP IDs) in the Path message so that the domain boundary node can identify which LSP segment or tunnel to use for the end-to-end LSP. Furthermore, this scheme may require communication to pre-establish the LSP segments.

These techniques may be directly applied, or may be applied with extensions, depending on specific diverse LSP setup schemes described below.

Note that in the schemes below, the paths of the working and recovery LSPs are not impacted by the confidentiality requirements.

6.1 Management Configuration

It is not possible to obtain or specify the full explicit route for either LSP at the head-end due to confidentiality restrictions. Therefore, this information cannot be provided to signaling for LSP setup.

Confidentially need not prevent the full computation of inter-domain paths and signaling of sufficient information to distinguish the paths. However the path information for each domain must be provided in a way that does not have meaning to other domains. Example mechanisms to preserve confidentiality as described above, include:

- Path key
- LSP segment

Details are for further study.

6.2 Head-end Path Computation (with multi-domain visibility)

This scheme is not applicable since multi-domain visibility violates confidentiality.

6.3 Per-Domain Path Computation

6.3.1 Sequential Path Computation

Assume the working LSP is established as described in [inter-domain-pd].

It is not possible to obtain the route of the working LSP from the RRO at the head-end due to confidentiality. In order to provide the path of the working LSP through each domain to the computation point responsible for computing the path of the protection LSP through each domain additional mechanisms are needed. Examples of such mechanisms are:

- Information identifying the working LSP is included in the Path message for the recovery LSP, and the domain boundary node consults the entity which computed the path of the working LSP (i.e., domain boundary node or PCE), so that the diverse path can be computed. When the entity which computed the path of the working LSP is the PCE, PCE needs to be temporarily stateful. An example of such information is the Association Object [e2e-recovery].

Details are for further study.

6.3.2 Simultaneous Path Computation

In this scheme the intention is to compute the path of the recovery LSP at the same time as the working LSP. In order to force the recovery LSP to follow the computed path as well as to preserve confidentiality, additional mechanisms are needed to communicate the computed recovery path from the path of the working LSP (where it was computed) to the recovery LSP. Examples of such mechanisms, that must continue to preserve confidentiality, are as follows.

- LSP segment: As described before. The LSP segment for the recovery LSP is established domain-by-domain as the working LSP setup progresses.

- Alternatively, information identifying the working LSP is included in the Path message for the recovery LSP, and the domain boundary node consults the entity which computed the path of the recovery LSP (i.e., domain boundary node or PCE), so as to obtain the path of the recovery LSP. This requires that the entity which computed the path of the recovery LSP is temporarily stateful. An example of such information is the Association Object [e2e-recovery].

Details are for further study.

6.4 Inter-domain Collaborate Path Computation

6.4.1 Sequential Path Computation

Assume working LSP is established as described in [brpc].

It is not possible to obtain RRO of working LSP (full route) at the head-end due to confidentiality.

o Path computation aspect

In order to obtain the path of the working LSP when computing the path of the recovery LSP, additional mechanisms are needed. Examples of such mechanisms are:

- Information identifying the working LSP is included in the PCEP message when requesting path computation of the recovery LSP should the PCE stateful (temporarily). An example of such information is the Association Object [e2e-recovery].

o Signaling aspect

The full route for the recovery LSP can not be returned to the head-end by PCE because it cannot be collected from the other PCEs owing to confidentiality restrictions. In order to force the recovery LSP to follow the path computed above, additional mechanisms are needed. Examples of such mechanisms are as described before:

- Path key
- LSP segment

Details are for further study.

6.4.2 Simultaneous Path Computation

It is not possible for PCE to return the full route of the working LSP and recovery LSP to the head-end due to confidentiality. In order to force the working and recovery LSPs to follow the paths computed, additional mechanisms are needed. Examples of such mechanisms are as described before:

- Path key: Use this for the working and recovery LSPs.

Note that the LSP segment approach in section 6 may not be applicable here since a path cannot be determined until BRPC procedures are completed.

Details are for further study.

7. Network Topology Specific Considerations

In some specific network topologies, diverse LSP setup schemes mentioned above could be drastically simplified.

For example, assume there are only three domains connected linearly, and the first and the last domain contain only a single node. Assume that we need to establish diverse LSPs from the node in the first domain to the node in the last domain. In such a scenario, no BRPC procedures are needed, because there is no need for path computation in the first and last domains.

8. Addressing Considerations

All of the above-mentioned schemes are applicable when a single address space is used across all domains.

However, there could be several cases where private addresses are used within some of the domains. This case is similar to the one with confidentiality, since the ERO and RRO are meaningless even if they are available. Details are for further study.

9. Note on SRLG Diversity

The above-mentioned schemes are applicable when the nodes and links in different domain belong to different SRLGs.

However, there could be several cases where the nodes and links in different domain belong to the same SRLG. That is, where SRLGs span domain boundaries. In such cases, in order to establish SRLG diverse LSPs, several considerations are needed, such as:

- Record of the SRLGs used by the working LSP
- Indication of a set of SRLGs to exclude in the computation of the recovery LSP's path.

Furthermore, SRLG IDs may be assigned independently in each domain, and might not have global meaning. In such a scenario, some mapping functions are necessary, similar to the mapping of other TE parameters mentioned in section 2.1.

Details are for further study.

10. Manageability Considerations

TBD

11. Security Considerations

TBD

12. References

12.1 Normative References

[inter-domain-fw] Farrel, A., et al., "A Framework for Inter-Domain MPLS Traffic Engineering", draft-ietf-ccamp-inter-domain-framework, work in progress.

[RFC4427] Mannie, E., Ed. and D. Papadimitriou, Ed., "Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4427, March 2006.

12.2 Informative References

[RFC4208] Swallow, G., et al., "Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", RFC 4208, October 2005.

[L1VPN-FW] Takeda, T., Editor "Framework and Requirements for Layer 1 Virtual Private Networks", draft-ietf-l1vpn-framework, work in progress.

[PCE-arch] A. Farrel, JP. Vasseur and J. Ash, "Path Computation Element (PCE) Architecture", draft-ietf-pce-architecture, work in progress.

[e2e-recovery] Lang, J., Rekhter, Y., and Papadimitriou, D. (Eds.), "RSVP-TE Extensions in support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS)-based Recovery", draft-ietf-ccamp-gmpls-recovery-e2e-signaling, work in progress.

[RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.

[seg-recovery] Berger, L., Bryskin, I., Papadimitriou, D., and Farrel, A., "GMPLS Based Segment Recovery", draft-ietf-ccamp-gmpls-segment-recovery, work in progress.

[RFC4105] Le Roux, J.-L. , Vasseur, J.-P., and J. Boyle, "Requirements for Inter-Area MPLS Traffic Engineering", RFC 4105, June 2005.

[RFC4216] Zhang, R., and Vasseur, J.-P., "MPLS Inter-Autonomous System (AS) Traffic Engineering (TE) Requirements", RFC 4216, November 2005

[inter-domain-rsvp] Ayyangar, A., Vasseur, JP. "Inter-domain MPLS Traffic Engineering - RSVP extensions", draft-ietf-ccamp-inter-domain-rsvp-te, work in progress.

[XRO] Lee et al., "Exclude Routes - Extension to RSVP-TE", draft-ietf-ccamp-rsvp-te-exclude-route, work in progress.

[inter-domain-pd] Vasseur JP., Ayyangar A., Zhang R., "A per-domain path computation method for computing Inter-domain Traffic Engineering Label Switched Path", draft-ietf-ccamp-inter-domain-pd-path-comp, work in progress.

[brpc] Vasseur, JP., Zhang, R., and Bitar, N., "A Backward Recursive PCE-based Computation (BRPC) procedure to compute shortest inter-domain Traffic Engineering Label Switched Path", draft-vasseur-ccamp-brpc, work in progress.

[PCEP] Vasseur, J., "Path Computation Element (PCE) communication Protocol (PCEP) - Version 1 -", draft-ietf-pce-pcep, work in progress.

[conf-segment] Bradford, R., Vasseur, JP., and Farrel, A., "Preserving Topology Confidentiality in Inter-Domain Path Computation and Signaling", draft-bradford-pce-path-key, work in progress.

[crankback] Farrel, A., et al., "Crankback Signaling Extensions for MPLS Signaling", draft-ietf-ccamp-crankback, work in progress.

[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.

[LSP-stitch] Ayyangar, A., and Vasseur, JP., "LSP Stitching

with Generalized MPLS TE", draft-ietf-ccamp-lsp-stitching, work in progress.

13. Acknowledgments

Authors would like to thank Eiji Oki, Ichiro Inoue and Kazuhiro Fujihara for valuable comments.

14. Author's Addresses

Tomonori Takeda
NTT Network Service Systems Laboratories, NTT Corporation
3-9-11, Midori-Cho
Musashino-Shi, Tokyo 180-8585 Japan
Email : takeda.tomonori@lab.ntt.co.jp

Yuichi Ikejiri
NTT Communications Corporation
Tokyo Opera City Tower 3-20-2 Nishi Shinjuku, Shinjuku-ku
Tokyo 163-1421, Japan
Email: y.ikejiri@ntt.com

Adrian Farrel
Old Dog Consulting
Email: adrian@olddog.co.uk

Jean-Philippe Vasseur
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough , MA - 01719
USA
Email: jpv@cisco.com

15. Intellectual Property Consideration

at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

16. Full Copyright Statement