Network Working Group   D. Fedyk
Internet Draft   Hewlett-Packard
Intended status: Standards Track   D. Beller
                                                           L. Levrau
                                                      Alcatel-Lucent
                                                       D. Ceccarelli
                                                            Ericsson
                                                            F. Zhang
                                                 Huawei Technologies
                                                           Y. Tochio
                                                            Fujitsu
                                                               X. Fu
                                                               ZTE

Expires: August 18, 2014   February 14, 2014

            UNI Extensions for Diversity and Latency Support
                draft-fedyk-ccamp-uni-extensions-04.txt

Status of this Memo

Copyright Notice

Abstract

This document builds on the GMPLS overlay model [RFC4208] and defines
extensions to the GMPLS User-Network Interface (UNI) to support route
diversity within the core network for sets of LSPs initiated by edge
nodes. A particular example where route diversity within the core
network is desired, are dual-homed edge nodes. The core network is
typically composed of multiple network domains and those multi-domain
diversity aspects that have an implication on the GMPLS UNI
extensions are discussed.

The document also defines GMPLS UNI extensions to deal with latency
requirements for edge node initiated LSPs.

This document uses a VPN model that is based on the same premise as
L1VPN framework [RFC4847] but may also be applied to other
technologies. The extensions are applicable both to VPN and non VPN
environments. These extensions move the UNI from basic connectivity
to enhanced mode connectivity by including additional constraints
while minimizing the exchange of CE to PE information. These
extensions are applicable to the overlay extension service model.
Route Diversity for customer LSPs are a common requirement applicable
to L1VPNs. The UNI mechanisms described in this document are L1VPN
compatible and can be applied to achieve diversity for sets of
customer LSPs.

The UNI extensions in support of latency constraints can also be
applied to the extended overlay service model in order for the
customer LSPs to meet certain latency requirements.

Table of Contents

1. Introduction

   This document builds on the GMPLS overlay model [RFC4208] and defines
   extensions to the GMPLS User-Network Interface (UNI) to support route
   diversity within the core network for sets of LSPs initiated by edge
   nodes. In the following, the term customer edge (CE) device is used
   synonymously for the term edge node (EN) as in [RFC4208].

   Moreover, the VPN terminology (CE and PE) [RFC4026] is used below
   when the core network is a VPN but is also applicable to UNI
   interfaces [RFC4208].

   This document uses a VPN model that is based on the same premise as
   L1VPN framework [RFC4847] but may also be applied to other

technologies. The extensions are applicable both to VPN and non VPN
environments. These extensions move the UNI from basic connectivity
to enhanced mode connectivity by including additional constraints
while minimizing the exchange of CE to PE information. These
extensions are applicable to the overlay extension service model.

The overlay model assumes a UNI interface between the edge nodes of
the respective transport domains. Route diversity for LSPs from
single homed CE and dual-home CEs is a common requirement in optical
transport networks. This document describes two signaling variations
that may be used for supporting LSP diversity within the overlay
extension service model considering dual-homing. Dual-homing is
typically used to avoid a single point of failure (UNI link, PE) or
if two disjoint connections are forming a protection group in the CE
device, e.g., 1+1 protection. While both methods are similar in that
they utilize common mechanisms in the PE network to achieve
diversity, they are distinguished according to whether the CE is
permitted to retrieve provider SRLG diversity information for an LSP
from a PE1 and pass it on to a PE2 (SRLG information is shared with
the CE), or whether a new attribute is used that allows the PE2 that
receives this attribute to derive the SRLG information for an LSP
based on the attribute value. Figure 1 below is depicting the
scenario.

The core network is typically composed of multiple network domains
(different providers, geographical separation, etc.) and some multi-
domain diversity aspects have implications on the GMPLS UNI
extensions defined in this document. It shall be noted that path
computation can be done in two different ways for each domain: GMPLS
supports distributed routing providing each node in the domain the
capability to do constraint-based path computation while the
utilization of the centralized path computation element (PCE)
approach is another option. The GMPLS UNI extensions defined in this
document are applicable to both path computation approaches and also
mixed scenarios are supported where some domains utilize the
distributed path computation approach while other domains are using a
PCE.

The extended overlay service model can support other extensions for
VPN signaling, for example, those related to latency. When requesting
diverse LSPs, latency may also be an additional requirement.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
document are to be interpreted as described in RFC-2119 [RFC2119].

In this document, these words will appear with that interpretation
only when in ALL CAPS. Lower case uses of these words are not to be
interpreted as carrying RFC-2119 significance.

3. Contributors

The Authors would like to thank Eve Varma and Sergio Belotti for
their review and contributions to this document.

4. LSP Diversity in the Overlay Extension Service Model

The L1VPN Framework [RFC4847] (Enhanced Mode) describes the overlay
extension service model, which builds upon the UNI Overlay [RFC4208]
serving as the interface between the CE edge node and the PE edge
node. In this service model, a CE receives a list of CE-PE TE link
addresses to which it can request a L1VPN connection (i.e.,
membership information) and may include additional information
concerning these TE links. This document further builds on the
overlay extension service model by adding shared constraint
information for path diversity in the optical transport network.
While the L1VPN for optical transport is an example specific VPN
technology the term VPN is used generically since the extensions can
apply to GMPLS UNIs and VPNs for other technologies.

Two signaling variations are outlined here that may be used for
supporting LSP diversity within the overlay extension service model
considering dual-homing. While both methods utilize common
mechanisms in the PE network to achieve diversity, they are
distinguished according to whether the CE is permitted to retrieve
provider SRLG diversity information for an LSP from a PE1 and pass it
on to a PE2 (SRLG information is shared with the CE or whether a new
attribute is used that allows the PE2 that receives this attribute to
derive the SRLG information for an LSP based on this attribute value.
The selection between these methods is governed by both PE-network
specific policies and approaches taken (i.e., in terms of how the
provider chooses to perform routing internal to their network).

The first method (see 4.1.1) assumes that provider Shared Resource
Link Group (SRLG) Identifier information is both available and
shareable (policy decision) with the CE. Since SRLG IDs can then be
used (passed transparently between PEs via the dual-homed CE) as
signaled information on a UNI message, a mechanism supporting LSP
diversity for the overlay extension service model can be provided via
straightforward signaling extensions.

The second method (see 3.1.2) assumes that provider SRLG IDs are
either not available or not shareable (based on provider network
operator policy) with the CE. For this case, a mechanism is provided

where information signaled to the PE on UNI messages does not require
shared knowledge of provider SRLG IDs to support LSP diversity for
the overlay extension model.

While both methods could be implemented in the same PE network, it is
likely that a GMPLS VPN CE network would use only one mechanism at a
time.

4.1. LSP diversity for dual-homed customer edge (CE) devices

Single-homed CE devices are connected to a single PE device via a
single UNI link (could be a bundle of parallel links which are
typically using the same fiber cable). This single UNI link may
constitute a single point of failure. Such a single point of failure
can be avoided when the CE device is connected to two PE devices via
two UNI interfaces as depicted for CE1 in Figure 1 below.

For the dual-homing case, it is possible to establish two connections
from the source CE device to the same destination CE device where one
connection is using one UNI link to, for example, PE1 and the other
connection is using the UNI link to PE2. In order to avoid single
points of failure within the provider network, it is necessary to
also ensure path (LSP) diversity within the provider network in order
to achieve end-to-end diversity for the two LSPs between the two CE
devices. This document describes how it is possible to enable such
path diversity to be achieved within the provider network (which is
subject to additional routing constraints). [RFC4202] defines SRLG
information that can be used to allow GMPLS to provide path diversity
in a GMPLS controlled transport network. As the two connections are
entering the provider network at different PE devices, the PE device
that receives the connection request for the second connection needs
to be capable of determining the additional path computation
constraints such that the path of the second LSP is disjoint with
respect to the already established first connection entering the
network at a different PE device. The methods described in this
document allow a PE device to determine the SRLG information for a
connection in the provider network that is entering the network on a
different PE device.

PE SRLG information can be used directly by a CE if the CE
understands the context, and the CE view is limited to its VPN
context. In this case, there is a dependency on the provider
information and there is a need to be able to query the SRLG in the
provider network.

It may, on the other hand, be preferable to avoid this dependency and
to decouple the SRLG identifier space used in the provider network
from the SRLG space used in the client network. This is possible with

both methods detailed below. Even for the method where provider SRLG
information is passing through the CE device (note the CE device does
not need to process and decode this information) the two SRLG
identifier spaces can remain fully decoupled and the operator of the
client network is free to assign SRLG identifiers from the client
SRLG identifier space to the CE to CE connection that is passing
through the provider network.

Referring to Figure 1, the UNI signaling mechanism must support at
least one of the two mechanisms described in this document for CE
dual homing to achieve LSP diversity in the provider network.

The described mechanisms can also be applied to a scenario where two
CE devices are connected to two different PE devices. In this case,
the additional information that is exchanged across the UNI
interfaces also needs to be exchanged between the two CE devices in
order to achieve the desired diversity in the provider network.

This information may be configured or exchanged by some automated
mechanism not described in this document.

In the dual-homing example, CE1 can locally correlate the LSP
requests. For the slightly more complicated example involving CE2 and
CE3, both requiring a path that shall be diverse to a connection
initiated by the other CE device, CE2 and CE3 need to have a common
view of the SRLG information to be signaled. In this document, we
detail the required diversity information and the signaling of this
diversity information; however, the means for distributing this
information within the PE domain or the CE domain is out of scope.

```
                      +---+     +---+
                      | P |....| P |
                      +---+     +---+
                     /               \
                  +-----+             +-----+     +---+
      +---+       | PE1 |             |     |----|   |
      |CE1|----|      |             |     |    |CE2|
      +---+\    +-----+             |     |----|   |
            \      |                | PE3 |     +---+
             \ +-----+             |     |
              \| PE2 |             |     |     +---+
               |     |             |     |----|CE3|
               +-----+             +-----+     +---+
                    \               /
                  +---+     +---+
                  | P |....| P |
                  +---+     +---+
```
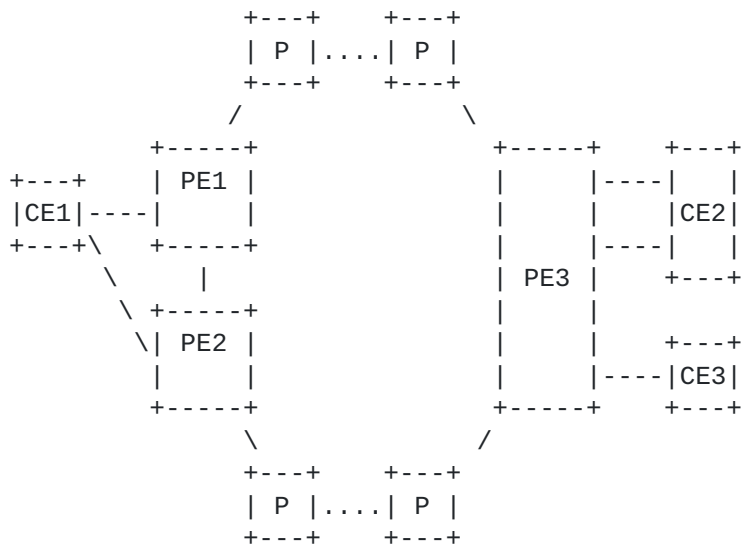
Figure 1 Overlay Reference Diagram

In an overlay model, the information exchanged between the CE and the PE is kept to a minimum.

How diversity is achieved, in terms of configuration, distribution and usage in each part of the transport networks should be kept independent and separate from how diversity is signaled at the UNI between the two transport networks.

Signaling parameters discussed in this document are:

o   SRLG information (see [RFC4202])

o   Path Affinity Set

4.1.1. **Exchanging SRLG information between the PEs via the CE device**

   SRLG information is defined in [RFC4202] and if the SRLG information
   of an LSP is known, it can be used to calculate a path for another
   LSP that is SRLG diverse with respect to an existing LSP. SRLG
   information is an unordered list of SRLGs. SRLG information is
   normally not shared between the transport network and the client
   network; i.e., not shared with the CEs of a VPN in the VPN context.
   However, this becomes more challenging when a CE is dual-homed. For
   example, CE1 in Figure 1 may have requested an LSP1 from CE1 to CE2
   via PE1 and PE3. CE1 could subsequently request an LSP2 to CE2 via
   PE2 and PE3 with the requirement that it should be maximally SRLG
   disjoint with respect to LSP1. Since PE2 does not have any
   information about LSP1, PE2 would need to know the SRLG information
   associated with LSP1. If CE1 could request the SRLG information of
   LSP1 from PE1, it could then transparently pass this information to
   PE2 as part of the LSP2 setup request, and PE2 would now be capable
   of calculating a path for LSP2 that is SRLG disjoint with respect to
   LSP1.

   The exchange of SRLG information is achieved on a per VPN LSP basis
   using the existing RSVP-TE signaling procedures. It can be exchanged
   in the PATH (exclusion information) or RESV message in the original
   request or it can be requested by the CE at any time the path is
   active.

   It shall be noted that SRLG information is an unordered list of SRLG
   identifiers and the encoding of SRLG information for RSVP signaling
   is already defined in [SRLG_info]. Even if SRLG information is known
   for several LSPs it is not possible for the CEs to derive the
   provider network topology from this information.

4.1.1.1. **Operational Procedures**

   Retrieving SRLG information from a PE for an existing LSP:

   When a dual-homed CE device intends to establish an LSP to the same
   destination CE device via another PE node, it can request the SRLG
   information for an already established LSP by setting the SRLG
   information flag in the LSP attributes sub-object of the RSVP PATH
   message (IANA to assign the new SRLG flag). As long as the SRLG
   information flag is set in the PATH message, the PE node inserts the

SRLG sub-object as defined in [SRLG_info] into the RSVP RESV message
that contains the current SRLG information for the LSP. If the
provider network's policy has been configured so as not to share SRLG
information with the client network, the SRLG sub-object is not
inserted in the RESV message even if the SRLG information flag was
set in the received PATH message. Note that the SRLG information is
expected to be always up-to-date.

Establishment of a new LSP with SRLG diversity constraints:

When a dual-homed CE device sends an LSP setup requests to a PE
device for a new LSP that is required to be SRLG diverse with respect
to an existing LSP that is entering the network via another PE
device, the CE device sets the SRLG diversity flag (note: IANA to
assign the new SRLG diversity flag) in the LSP attributes sub-object
of the PATH message that initiates the setup of this new LSP. When
the PE device receives this request it calculates a path to the given
destination and uses the received SRLG information as path
computation constraints.

4.1.1.2. Error Handling Procedures

When the CE device receives a RSVP PATH message with the SRLG
information flag set and if the provider's network policy does not
permit sharing of SRLG information, the PE device shall notify the CE
device by sending a RSVP PathErr with a Notify error code (error code
to be defined) "Retrieval of SRLG information not permitted". As
described above, the PE device must not include the SRLG sub-object
with the SRLG information for the LSP in the RSVP RESV message.

If the PE device receives a RSVP PATH message for a new LSP with the
SRLG diversity flag set and SRLG information in the SRLG sub-object,
the PE device tries to calculate a route to the given destination
that is SRLG diverse with respect to the provided SRLG information.
If no route can be found, a RSVP PathErr message with an error code
(error code to be defined) "No SRLG diverse route available toward
destination".

If the PE device receives a RSVP PATH message for a new LSP with the
SRLG diversity flag set and SRLG information in the SRLG sub-object
and if the PE device does not support the SRLG sub-object, the PE
device shall send a PathErr message to the CE device, indicating an
"Unknown object class".

Further error handling cases will be added in the next revision of

this document.

4.1.2. Using Path Affinity Set Extension

The Path Affinity Set (PAS) is used to signal diversity in a pure CE
context by abstracting SRLG information. There are two types of
diversity information in the PAS. The first type of information is a
single PAS identifier. The Second part is the optional PATH
information, in the form of Source and Destination addresses of an
exclude path or set of paths that MAY be specified. The motive behind
the PAS information is to have as little exchange of diversity
information as possible between the VPN CE and PE elements.

Rather than a detailed CE or PE SRLG list, the Path Affinity Set
contains an abstract SRLG identifier that associates the given path
as diverse. Logically the identifier is in a VPN context and
therefore only unique with respect to a particular VPN.

How the CE determines the PAS identifier is a local matter for the CE
administrator. A CE may signal the PAS identifier as a diversity
object in the PATH message. This identifier is a suggested identifier
and may be overridden by a PE under some conditions.

For example, a PAS identifier can be used with no prior exchange of
PAS information between the CE and the PE. Upon reception of the PAS
identifier information the PE can infer the CEs requirements. The
actual PAS identifier used will be returned in the RESV message.
Optionally an empty PAS identifier allows the PE to pick the PAS
identifier.

Similar to the section 4.1.1 on SRLG information, a PE can return PAS
identifier as the response to a Query allowing flexibility.

A PE interprets the specific PAS identifier, for example, "123" as
meaning to exclude the PE SRLG information (or equivalent) that has
been allocated by LSPs associated with this Path Affinity Set
identifier "123", for any LSPs associated with the resources assigned
to the VPN. For example, if a Path exists for the LSP with the
identifier "123", the PE would use local knowledge of the PE SRLGs
associated with the "123" LSPs and exclude those SRLGs in the path
request. In other words, two LSPs that need to be diverse both
signal "123" and the PEs interpret this as meaning not to use shared
resources. Alternatively, a PE could use the PAS identifier to
select from already established LSPs. Once the path is established it
becomes the "123" identifier or optionally another PAS identifier for
that VPN that replaces "123".

The optional PAS Source and Destination Address tuple represents one
or more source addresses and destination addresses associated with
the CE Path Affinity Set identifier. These associated address tuples
represent paths that use resources that should be excluded for the
establishment of the current LSP. The address tuple information
gives both finer grain details on the path diversity request and
serves as an alternative identifier in the case when the PAS
identifier is not known by the PE. The address tuples used in
signaling is within a CE context and its interpretation is local to a
PE that receives a Path request from a CE. The PE can use the address
information to relate to PE Addresses and PE SRLG information. When
a PE satisfies a connection setup for a (SRLG) diverse signaled path,
the PE may optionally record the PE SRLG information for that
connection in terms of PE based parameters and associate that with
the CE addresses in the Path message.

Specifically for L1VPNs, Port Information table (PIT) [RFC5251] can
be leveraged to translate between CE based addresses and PE based
addresses. The Path Affinity Set and associated PE addresses with PE
SRLG information can be distributed via the IGP in the provider
transport network (or by other means such as configuration); they can
be utilized by other PEs when other CE Paths are setup that would
require path/connection diversity. This information is distributed on
a VPN basis and contains a PAS identifier, PE addresses and SRLG
information.

If diversity is not signaled, the assumption is that no diversity is
required and the Provider network is free to route the LSP to
optimize traffic. No Path affinity set information needs to be
recorded for these LSPs. If a diversity object is included in the
connection request, the PE in the Provider Network should be able to
look-up the existing Provider SRLG information from the provider
network and choose an LSP that is maximally diverse from other LSPs.

The mechanisms to achieve this are outside the scope of this
document.

A new VPN Diverse LSP LABEL object is specified:


```
    0                   1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |       Length                  |    Type (TBA) |0| C-type (TBA)|
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+

                        1                   2                   3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1

   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   | ADDR Length   |Number of PAS  |D|            reserved         |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                Path Affinity Set identifier                  |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                 Source Address (variable)                    |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                Destination Address (variable)                |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```
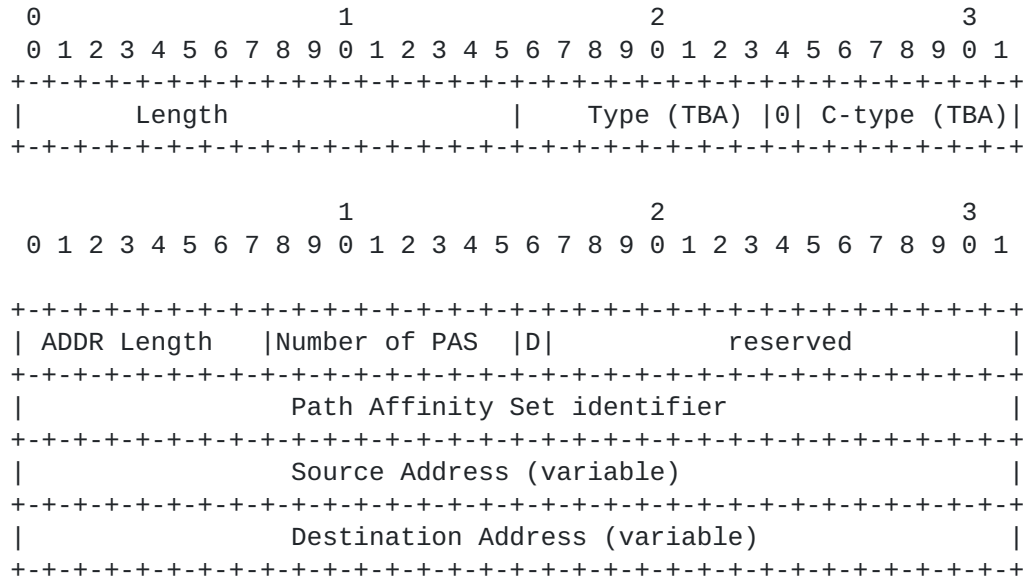
                    Figure 2 Diverse LSP information

1. The Address Length field (8 bits) is the number of bytes for both
   the source address and destination address. The address may be in
   any format from 1 to 32 bytes but the key point is the customers
   can maintain their existing addresses. A value of zero indicates
   there are no addresses included.

2. The Number of Path Affinity (8 bits) sets is included in the
   object. This is typically 1. Addition of other sets is for further
   study.

3. The Path affinity Set identifier (4 bytes) is a single number that
   represents a summarized SRLG for this path. Paths with that same
   Path Affinity set should be set up with diverse paths and
   associated with the path affinity set. A value of all zeros
   allows the PE to pick a PAS identifier to return. A PAS
   identifier of an established path may be different than the
   requested path identifier.

4. The diversity Bit (D) (one Bit) indicates if the diversity must be
   satisfied when set as a one. If a PE finds an established path
   with a Path Affinity set matching the signaled Path Affinity Set
   or the signaled Address tuple it should attempt find a diverse
   path.

5. The Diverse Path Source address/destination address tuple is that
   of an established LSP in the PE network that belongs to the same
   Path Affinity Set identifier. If the path for these addresses is
   not established or cannot be determined by the PE edge processing
   the PATH request then the path is established only with the Path
   Affinity identifier. If the path(s) for these address tuples are
   known by the PE the PE uses the SRLG information associated with
   these addresses. If in any case a diverse path cannot be setup
   then the Diverse bit controls whether a path is established
   anyway. The PE must use the PIT to translate CE Addresses into
   provider addresses when correlating with provider SRLG
   information. How SRLG information and network address tuples are
   distributed is for future study.

4.1.2.1. Operational Procedures

When a CE constructs a PATH message it may optionally specify and
insert a Path Affinity Set in the PATH message. This Path Affinity
Set may optionally include the address of an LSP that that could
belong to the same Path Affinity Set. The Path Affinity Set
identifier is a value (0 through 2**32-255) that is independent of
the mechanism the CE or the PE use for diversity. The Path Affinity
Set is a single identifier that can be used to request diversity and
associate diversity.

When processing a CE PATH message in a VPN Overlay, the PE first
looks up the PE based addresses in the Provider Index Table (PIT). If
the Path Affinity Set is included in the PATH message, the PE must
look up the SRLG information (or equivalent) in the PE network that
has been allocated by LSPs associated with a Path Affinity Set and
exclude those resources from the path computation for this LSP if it
is a new path. The PE may alternatively choose from an existing path
with a disjoint set of resources. If a path that is disjoint cannot
be found, the value of the PAS diversity bit determines whether a
path should be setup anyway. If the PAS diversity bit is clear, one
can still attempt to setup the LSP. A PE should still attempt to
minimize shared resources but that is an implementation issue, and is
outside the scope of this document.

Optionally the CE may use a value of all zeros in the PAS identifier
allowing the PE to select an appropriate PAS identifier. Also the PE
may to override the PAS identifier allowing the PE to re-assign the
identifier if required. A CE should not assume that the PAS
identifier used for setup is the actual PAS identifier.

4.1.2.2. Error Handling Procedures

The PAS object must be understood by the PE device. Otherwise, the CE
should not use the PAS object. Path Message processing of the PAS
object SHOULD follow CTYPE 0. An Error code of IANA (TBD) indicates
that the PAS object is not understood.

When a PAS identifier is not recognized by a PE it must assume this
LSP defines that PAS identifier however the PE may override PAS
identifier under certain conditions.

If the identifier is recognized but the Source Address-Destination
address pair(s) are not recognized, this LSP must be set up using the
PAS identifier only.

If the identifier is recognized and the Source Address-Destination
address pair(s) are also recognized, then the PE SHOULD use the PE
SRLG information associated with the LSPs identified by the address
pairs to select a disjoint path.

The Following are the additional error codes:

1. Route Blocked by Exclude Route Value IANA (TBA).

4.1.2.3. Distribution of the Path Affinity Set Information

Information about SRLG is already available in the IGP TE database. A
PE network can be designed to have additional opaque records for
Provider paths that distribute PE paths and SRLG on a VPN basis. When
a PE path is setup, the following information allows a PE to lookup
the PE diversity information:

o L1 VPN Identifier 8 bytes

o Path Affinity Set Identifier

o Source PE Address

o Destination PE Address

o List of PE SRLG (variable)

The source PE address and destination PE address are the same
addresses in the VPN PIT and correspond to the respective CE address
identifiers.

   Note that all of the information is local to the PE context and is
   not shared with the CE. The VPN Identifier is associated with a CE.
   The only value that is signaled from the CE is the Path Affinity Set
   and optionally the addresses of an existing LSP. The PE stores source
   and destination PE addresses of the LSP in their native format along
   with the SRLG information. This information is internal to the PE
   network and is always known.

   PE paths may be setup on demand or they may be pre-established. When
   paths are pre-established, the Path Affinity Set is set to unassigned
   0x0000 and is ignored. When a CE uses a pre-established path the PE
   may set the Path SRLG Path Affinity Set value if the CE signals one
   otherwise the Path Affinity Set remains unassigned 0x0000.

4.2. Multi-domain LSP Diversity Aspects for Dual-homed CE Devices

   The two mechanisms described above to achieve LSP diversity for
   dual-homed CE devices can be applied to single-domain provider
   networks as well as multi-domain provider networks. This section
   addresses multi-domain aspects including both single provider multi-
   domain networks and multi-provider networks where the subdivision
   into multiple domains is obvious due to the organizational boundaries
   between different providers. Specifically, when multiple providers
   are involved, SRLG identifiers as well as PAS identifiers must be
   administrable independently for each provider network.

   For the single provider multi-domain case, there are two
   possibilities how SRLG or PAS identifiers can be handled:

   o Subdividing the identifier space into ranges assigned to domains

   o Scoping the identifiers to domains

4.2.1 Subdividing Identifier Spaces into Ranges

   Subdividing the identifier space into disjoint ranges and assigning
   the different ranges to the different domain is one possibility to
   apply the LSP diversity mechanisms defined in this document to a
   multi-domain environment. This does not require additional protocol
   extensions. Caution is, however, required when the identifiers are
   assigned. They must be selected strictly from the identifier range
   that has been assigned to the specific domain. From a network
   operations perspective, this can be an option for a single provider
   multi-domain network while it may be less applicable to multi-
   provider networks where minimal dependency is desired.

4.2.2 Scoping Identifier Spaces to Domains

[DRAFT DOMAIN SUBOBJECTS] defines new RSVP-TE domain sub-objects for
the purpose of identifying domains. Domain sub-objects can be used to
scope SRLG or PAS identifiers to a specific domain. With this
extension, the full SRLG or PAS identifier space can be used within
each domain. When a new multi-domain LSP shall be established, the
diversity constraints can be signaled in the form of a sequence of a
scoping domain sub-object followed by the list of SRLGs (SRLG sub-
object) or the PAS sub-object, e.g.: [domain_sub-object(Dn),
SRLG_sub-object(Dn)] for domain Dn.

4.2.3. Multi-domain Diversity Aspects in Case Domains Utilize a PCE

Typically, the core network is composed of multiple network domains
(different providers, geographical separation, etc.) and some multi-
domain diversity aspects have implications on the GMPLS UNI
extensions defined in this document.

For GMPLS controlled networks, two options are defined how path
computation can be done:

o Distributed path computation, i.e., each node is capable to
perform constraint-based path computation

o Centralized path computation utilizing PCE as defined in [RFC4655]

The GMPLS UNI extensions defined in this document shall be applicable
to both path computation approaches and also mixed scenarios shall be
supported where some domains utilize the distributed path computation
approach while other domains are using a PCE.

In case a network domain uses a PCE, path information for all LSPs
crossing the domain can be stored in the PCE's database and [DRAFT
PATH KEY] defines a mechanism how a LSP diversity constraint can be
signaled in the RSVP-TE eXclude Route Object (XRO) using a unique
path key encoded in a path key sub-object. Further details can be
found in [DRAFT PATH KEY].

If the scoping approach as defined in section 4.2.2 above is applied,
the diversity constraint for an LSP can be signaled in the form of a
sequence of a domain sub-object followed by a path key sub-object
and the path key sub-object itself contains the PK-owner-ID that
tells the ingress node of a domain receiving the diversity constraint
which PCE instance it has to consult.

For mixed scenarios, where some domains are using the distributed
path computation approach while the other domains are utilizing a
PCE, the LSP diversity constraint can be signaled in the form of a

sequence of a scoping domain sub-object followed by the list of SRLGs
(SRLG sub-object) or the PAS sub-object (distributed path
computation) or the path key sub-object for domains using a PCE.
Hence the diversity constraint for a domain Dn has the following
form:

[domain_sub-object(Dn), SRLG_sub-object(Dn) | PAS_subobject(Dn) |
PK_sub-object(Dn)]

5. Latency Signaling Extensions

Some network applications are sensitive to latency (sometimes also
called delay) while other applications are sensitive to latency
variation (sometimes also called delay variation). Specifically, real
time applications typically do have certain latency requirements. It
shall be noted that latency variation is typically not an issue for
TDM networks including the WDM layer. For these technologies the
latency is constant and there is no latency variation added. Latency
variation is typically caused in packet networks or when packet based
services are encapsulated into a constant bit rate server layer
signal, which requires buffering of the arriving packets that may
arrive in bursts. An example is an Ethernet VLAN service that is
mapped into a constant bit rate server layer such as an ODUk or
ODUflex OTN signal.

The GMPLS UNI as defined in [RFC4208] does not support latency as a
signaling parameter that would allow a CE device to signal to the PE
device that latency and/or latency variation constraints need to be
met when a path is calculated for the requested LSP. The path
computation function does typically calculate a route to the given
destination that has the least TE metric (least cost routing).
However, if a CE device requests an LSP via the UNI interface for an
application that is sensitive to latency/latency variation, it should
be possible to signal to the PE device that the objective function
should rather take latency into account instead of the TE metric.

In order to support latency/latency variation as path computation
constraint, the network has to support latency/latency variation as
TE metric extension as defined in [DRAFT OSPF TE METRIC EXT] - note
that [DRAFT OSPF TE METRIC EXT] is using the terms delay/delay
variation instead of latency/latency variation.

A latency requirement can be added to signaling in the form of a
constraint [DRAFT OBJECTIVE FUNCTION]. The constraint can take the
form of:

o Minimal latency

o Maximum acceptable latency (upper bound)

o Minimal latency variation

o Maximum acceptable latency variation (upper bound), if applicable

While some systems may be able to compute routes based on delay
metrics it is usual that minimizing the accumulated TE link metric
(link cost) or the number of hops subject to bandwidth reservation
are satisfied as the object function and delay is not considered.
When considering diversity latency falls after diversity constraints
have been satisfied.

Recording the latency of existing paths [DRAFT TE METRIC RECORD] to
ensure they meet a maximum acceptable latency can be utilized to
ensure latency constraint is met.

When a low latency path is required, the minimize latency subject to
other constraints criteria should be signaled. A CE device can use
the recorded latency to ensure that the maximum acceptable latency
has been met.

5.1. RSVP-TE Extensions

At the UNI, the RSVP-TE extensions as defined in [DRAFT OBJECTIVE
FUNCTION] SHALL be used for signaling the PE device whether a path
with minimal latency is requested or whether certain latency/latency
variation upper bound constraints shall be met for the end-to-end
connection, i.e., from the source CE device to the destination CE
device. The following objective function (OF) code point SHALL be
used in the OF sub-object of the ERO to indicate that latency/latency
variation constraints SHALL be taken into account when the path
computation function that is invoked by the PE node that expands the
route from the PE device to the destination CE device:

o OF code value 8 (to be assigned by IANA) is for the Minimum
   Latency Path (MLP) OF

o OF code value 9 (to be assigned by IANA) is for Minimum Latency
   Variation Path (MLVP) OF

Additionally, an optional OF metric-bound sub-object MAY be carried
within an ERO object of the RSVP-TE Path message. The two metric-
bound sub-objects defined in [DRAFT OBJECTIVE FUNCTION] that are
corresponding to the two OFs above are:

o metric bound sub-object of Type T=4: Cumulative Latency

o metric bound sub-object of Type T=5: Cumulative Latency Variation

The metric-bound indicates an upper bound for the path metric that
MUST NOT be exceeded for the ERO expending node to consider the
computed path as acceptable. It shall be noted that the metric bound
included in the RSVP-TE Path message at the UNI has end-to-end
significance, which means that the upper bound metric constraint MUST
be met for the path from the source CE device to the destination CE
device.

5.2. Operational Procedures

The processing rules as defined in [DRAFT OBJECTIVE FUNCTION] for the
OF sub-object and the optional OF metric-bound sub-object SHALL be
applied at the ingress PE device when the source CE device requests
an LSP (It shall be noted that [DRAFT OBJECTIVE FUNCTION] has a wider
scope and may also apply to inter-domain interfaces, i.e., when the
provider network is composed of multiple separate domains.).

5.3. Error Handling Procedures

The error handling rules as defined in [DRAFT OBJECTIVE FUNCTION] for
the OF sub-object and the optional OF metric-bound sub-object SHALL
be applied.

6. Security Considerations

Security for L1VPNs is covered in [RFC4847], [RFC5251] and [RFC5253].
In this document, the model follows a generic GMPLS VPN based on the
L1VPN control plane model where CE addresses are completely distinct
from the PE addresses.

The use of a private network assumes that entities outside the
network cannot spoof or modify control plane communications between
CE and PE. Furthermore, all entities in the private network are
assumed to be trusted. Thus, no security mechanisms are required by
the protocol exchanges described in this document.

However, an operator that is concerned about the security of their
private control plane network may use the authentication and
integrity functions available in RSVP-TE [RFC3473] or utilize IPsec
([RFC4301], [RFC4302], [RFC4835], [RFC5996], and [RFC6071]) for the
point-to-point signaling between PE and CE. See [RFC5920] for a full
discussion of the security options available for the GMPLS control
plane.

7. IANA Considerations

TBD

8. References

8.1. Normative References

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

[RFC4202] Kompella, K., Rekhter, Y., "Routing Extensions in Support
          of Generalized Multi-Protocol Label Switching (GMPLS)", RFC
          4202, October 2005.

[RFC4208] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter,
          "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.

[RFC4655] Farrel, A., Vasseur, J.-P., Ash, J., "A Path Computation
          Element (PCE)-Based Architecture", RFC4655, August 2006.

[RFC5251] Fedyk, D., Rekhter, Y., Editors "Layer 1 VPN Basic Mode",
          RFC 5251, July 2008.

[SRLG_info] Zhang, F., Gonzalez de Dios, O., Li, D., Margaria, C.,
          Hartley, M., Ali, Z., "RSVP-TE Extensions for Collecting
          SRLG Information", draft-ietf-ccamp-rsvp-te-srlg-collect-
          04.txt, February 2014.

[DRAFT OBJECTIVE FUNCTION] Ali, Z., Swallow, G., Filsfils, C., Fang,
          L., Kumaki, K., Kunze, R., Ceccarelli, D., Zhang, X.,
          "Resource ReserVation Protocol - Traffic Engineering
          (RSVP-TE) extension for signaling Objective Function and
          Metric Bound", draft-ali-ccamp-rc-objective-function-
          metric-bound-04.txt, October 2013.

[DRAFT DOMAIN SUBOBJECTS] Dhody, D., Palle, U., Kondreddy, V.,
          Casellas, R., "Domain Subobjects for Resource ReserVation
          Protocol - Traffic Engineering (RSVP-TE)", draft-ietf-
          ccamp-rsvp-te-domain-subobjects-01.txt, January 2014.

[DRAFT PATH KEY] Zhang, X., Zhang, F., Gonzalez de Dios, O., Bryskin,
          I., Dhody, D., "Extensions to Resource ReSerVation
          Protocol-Traffic Engineering (RSVP-TE) to Support Route
          Exclusion Using Path Key Subobject", draft-zhang-ccamp-
          route-exclusion-pathkey-01.txt, February 2014

8.2. Informative References

[RFC4026] Andersson, L. and T. Madsen, "Provider Provisioned Virtual
          Private Network (VPN) Terminology", RFC 4026, March 2005.

[RFC6071] Frankel, S. and S. Krishnan, "IP Security (IPsec) and
          Internet Key Exchange (IKE) Document Roadmap", RFC 6071,
          February 2011.

[RFC3473] Berger, L. (editor), "Generalized MPLS Signaling - RSVP-TE
          Extensions", RFC 3473, January 2003.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the
          Internet Protocol", RFC 4301, December 2005.

[RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December
          2005.

[RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet
          Key Exchange Protocol Version 2 (IKEv2)", RFC 5996,
          September 2010.

[RFC4835] Manral, V., "Cryptographic Algorithm Implementation
          Requirements for Encapsulating Security Payload (ESP) and
          Authentication Header (AH)", RFC 4835, April 2007.

[RFC4847] Takeda, T., Editor "Framework and Requirements for Layer
          Virtual Private Networks", RFC 4847, April 2007.

   [RFC5253] Takeda, T., Ed., "Applicability Statement for Layer 1
             Virtual Private Network (L1VPN) Basic Mode", RFC 5253, July
             2008.

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

   [DRAFT TE METRIC RECORD] Ali, Z., Swallow, G., Filsfils, C., Hartley,
             M., Kumaki, K., Kunze, R., "Resource ReserVation Protocol-
             Traffic Engineering (RSVP-TE) extension for recording TE
             Metric of a Label Switched Path", draft-ietf-ccamp-te-
             metric-recording-02.txt, July 2013.

   [DRAFT OSPF TE METRIC EXT] Giacalone, S., Ward, D., Drake, J., Atlas,
             A., Previdi, S., "OSPF Traffic Engineering (TE) Metric
             Extensions", draft-ietf-ospf-te-metric-extensions-05.txt,
             December 2013.

Authors' Addresses

   Don Fedyk
   Hewlett-Packard
   153 Tayor Street
   Littleton, MA, 01460
   Email: don.fedyk@hp.com

   Dieter Beller
   Alcatel-Lucent
   Email: Dieter.Beller@alcatel-lucent.com

   Lieven Levrau
   Alcatel-Lucent

   Email: Lieven.Levrau@alcatel-lucent.com

   Daniele Ceccarelli
   Ericsson
   Email: Daniele.Ceccarelli@ericsson.com

   Fatai Zhang
   Huawei Technologies
   Email: zhangfatai@huawei.com

   Yuji Tochio
   Fujitsu
   Email: tochio@jp.fujitsu.com

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