

(43) International Publication Date
16 January 2014 (16.01.2014)(51) International Patent Classification:
H04L 12/735 (2013.01)(21) International Application Number:
PCT/EP2013/062501(22) International Filing Date:
17 June 2013 (17.06.2013)

(25) Filing Language: English

(26) Publication Language: English


(30) Priority Data:
12305817.4 9 July 2012 (09.07.2012) EP(71) Applicant: **ALCATEL LUCENT** [FR/FR]; 3, avenue
Octave Gréard, F-75007 Paris (FR).(72) Inventors: **BELLER, Dieter**; Hoffmannstr. 14/3, 70825
Kornthal (DE). **FEDYK, Donald**; 220 Hayden Road, Gro-
ton, MA, 01450 (US).(74) Agent: **URLICHS, Stefan**; Alcatel-Lucent Deutschland
AG, Intellectual Property Business Group, Lorenzstr. 10,
70435 Stuttgart (DE).(81) Designated States (*unless otherwise indicated, for every
kind of national protection available*): AE, AG, AL, AM,
AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,
BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM,DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,
HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KN, KP, KR,
KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME,
MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ,
OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC,
SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN,
TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.(84) Designated States (*unless otherwise indicated, for every
kind of regional protection available*): ARIPO (BW, GH,
GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ,
UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,
TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK,
EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV,
MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM,
TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).**Declarations under Rule 4.17:**

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *of inventorship (Rule 4.17(iv))*

Published:

- *with international search report (Art. 21(3))*

(54) Title: METHOD AND RELATED APPARATUS FOR ESTABLISHING LINK-DIVERSE TRAFFIC PATHS IN A TELECOMMUNICATIONS NETWORK

**Fig. 2**

(57) **Abstract:** In order to achieve path diversity for dual-homed UN I-clients connected to a GMPLS control plane enabled transport network that is operated in an overlay mode, the overlay extension service model is enhanced by adding shared constraint information for path diversity in the optical transport network. In particular, within the provider network, shared constraint information of a first traffic path (LSP1) is determined and a data element indicative of the shared constraint information is returned by a first provider edge node (PE1) to a customer edge device (CE1) via a User-Network Interface. When the customer edge device (CE1) requests a second traffic path (LSP2) through the provider network to be established from a second provider edge node (PE2) and to be disjoint from said first traffic path (LSP1), the customer edge device forwards the data element to the second provider edge node (PE2) to enable path calculation of the second traffic path using the shared constraint information as exclusion list.

5

10 **METHOD AND RELATED APPARATUS FOR ESTABLISHING LINK-DIVERSE
TRAFFIC PATHS IN A TELECOMMUNICATIONS NETWORK**

15 **Field of the Invention**

The present invention relates to the field of telecommunications and more particularly to a method and related apparatus for establishing link-diverse traffic paths in a telecommunications network.

20 **Background of the Invention**

In classical layer 1 transport networks employing time-division multiplexing (TDM) and/or wavelength division multiplexing (WDM), such as the Optical Transport Network (OTN) or the Synchronous Digital Hierarchy (SDH),
25 techniques have been established to support dynamic resource allocation and setup of traffic paths. A transport network, which allows traffic paths to be set up through a switched network automatically is termed Automatically Switched Transport Network (ASTN) or Automatically Switched Optical Network (ASON). The WDM layer is often also called layer
30 0.

Traditionally, creating traffic paths through a series of Network Elements has involved configuration of individual cross-connections on each network element. ASTN/ASON allows the user to specify the start point, end point,

and bandwidth required, and the ASTN/ASON agent on the network elements will allocate the path through the network, provision the traffic path, set up cross-connects, and allocate bandwidth from the paths for the user requested service. The actual path that the traffic will take through
5 the network is not specified by the user.

A protocol suite known as GMPLS (Generalized Multiprotocol Label Switching) has been developed to dynamically provision resources and to provide network survivability using protection and restoration techniques.
10 GMPLS includes three main protocols: a signaling protocol known as the Resource Reservation Protocol with Traffic Engineering extensions (RSVP-TE), a routing protocol known as the Open Shortest Path First with Traffic Engineering extensions (OSPF-TE), and the Link Management Protocol (LMP).

15 The underlying architecture of an ASTN/ASON assumes a forwarding (or data) plane and a separated control plane, where each of the network layers can use physically diverse data or forwarding planes. The focus of GMPLS is on the control plane.

20 Furthermore, the concept of Virtual Private Networks (VPNs) has been extended to layer 1 networks, recently. A Layer 1 VPN (L1VPN) is a service offered by a core layer 1 network to provide layer 1 connectivity between two or more customer sites, and where the customer has some control over
25 the establishment and type of the connectivity (see IETF RFC 4847).

GMPLS defines both routing and signaling protocols for the creation of Label Switched Paths (LSPs). In a peer model, edge nodes support both a routing and a signaling protocol. The protocol interactions between an edge node
30 and a core node are the same as between two core nodes. In the overlay model, however, the core nodes act more as a closed system. The edge nodes do not participate in the routing protocol instance that runs among the core nodes; in particular, the edge nodes are unaware of the topology of the core nodes (see IETF RFC 4208).

In the overlay model, customer edge (CE) devices are connected over a User-Network Interface (UNI) to provider edge (PE) devices. PE devices are connected to other provider (P) devices via a Network-Network Interface (NNI).

Current UNIs include features to facilitate requests for end-to-end (i.e. CE to CE) services that include the specification of constraints such as bandwidth requirements, protection needs, and (of course) destinations. Current NNIs on the other hand include features to exchange routing information, as well as to facilitate requests for end-to-end services.

If one CE device is connected to two PE devices, this configuration is referred to as dual-homing. Dual-homing is typically used to avoid a single point of failure (e.g. UNI link or PE) or if two disjoint connections will form a protection group. Route diversity for LSPs from dual-homed CEs is a common requirement in optical transport networks.

Summary of the Invention

A problem of the present invention is to achieve path diversity for dual-homed UNI-clients connected to a GMPLS control plane enabled transport network that is operated in an overlay mode.

EP2267954A1 describes a solution where network resources are assigned to different sets of resources in such a way that diversity is achieved if a first connection uses resources from a first resource set, only, while a second connection uses only resources from a second resource set. The sets must be constructed in such a way that path diversity is achieved provided that the two connections use different resource sets. This solution is somehow restrictive. Moreover, it requires more network resources and adds additional penalties to the network planning process as well as to the network operation.

US 2009/0103442 A1 describes constraint-based routing in multi-domain networks. For a connection path within a domain, a set of shared risk groups (SRGs) associated with resources used by the path is collected and stored within a shared risk group management device associated with that domain. A shared risk identification code is assigned to the set of shared risk groups. Risk information which contains the shared risk identification code and a shared risk groups management device identifier is communicated to the neighboring domain. This allows the neighboring domain to request a path to be routed SRG-disjoint by using the risk information to exclude resources that would share a same SRG.

The below described embodiments introduce signaling variations that may be used for supporting LSP diversity within the overlay extension service model considering dual-homing.

The L1VPN Framework in RFC4847 describes the overlay extension service model, which builds upon the UNI Overlay (RFC4208) serving as the interface between the CE device and the PE 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. The described embodiments further build on the overlay extension service model by adding shared constraint information for path diversity in the optical transport network. This allows a different network entry point to calculate a path that is divers with respect to an existing path through the network that has a different network entry point.

In particular, within the provider network, shared constraint information of a first traffic path is determined and a data element indicative of the shared constraint information is returned by a first provider edge node to a customer edge device via a User-Network Interface. The shared constraint information contains Shared Risk Link Group identifiers of all provider network links used by the first traffic path and the data element contains an unordered list of Shared Risk Link Group identifiers of the first traffic

- path. When the customer edge device requests a second traffic path through the provider network to be established from a second provider edge node and to be disjoint from said first traffic path, the customer edge device forwards the data element to the second provider edge node to
- 5 enable path calculation of the second traffic path using the shared constraint information as exclusion list.

Brief Description of the Drawings

- 10 Preferred embodiments of the present invention will now be described with reference to the accompanying drawings in which
- figure 1 shows an example network topology;
- figure 2 shows a first signaling variant to exchange shared constraint information;
- 15 figure 3 shows a new L1VPN Diverse LSP LABEL object used to signal shared constraint information to and from a customer edge device;
- figure 4 shows second signaling variant to exchange shared constraint information;
- 20 figure 5 shows a block diagram of a provider edge node; and
- figure 6 shows a block diagram of a customer edge device.

Detailed Description of the Invention

- An simplified network example is shown in figure 1. A first customer edge device CE1 is connected in a dual-homed configuration to a first and a second provide edge node PE1, PE2. A second customer edge device CE2 is connected to provider edge nodes PE3, PE4. The provider network contains provider core nodes P1-P4. Provider edge nodes PE1-PE4 and provider core nodes P1-P4 are interconnected in a mesh topology, which is not shown for the sake of simplification. Connections are established over label switched paths LSP1, LSP2 between CE1 and CE2 via PE1, P1, P2, and PE3 and via PE2, P3, P4, and PE4. Label Switched Paths LSP1 and LSP2 are link diverse, so that no single point of failure exists for the two connections.
- The interfaces between customer edge devices CE1, CE2 and provider edge nodes PE1, PE2, PE3, PE4, respectively, are User-Network Interfaces UNI, while interfaces between provider edge nodes PE1-PE4 and provider core nodes P1-P4 are Network-Network Interfaces NNI.
- On the NNI, all provider nodes PE1-PE4, P1-4 exchange routing protocol messages using a routing protocol such as OSPF-TE to advertise topology and link state information, which is used to populate local routing databases in each provider node PE1-PE4, P1-4. Such routing messages are broadcasted to all other provider nodes, so that by listening to the routing messages, all provider nodes share the same topology and link state information about the entire provider network. The term routing database as used herein refers to the sum of all routing information stored in a PE node and can include an adjacency database, a link state database and a Traffic Engineering database.
- On the UNI, however, no such routing information is exchanged, so that the CE devices CE1, CE2 do not receive any topology and link state information about the provider network. This is intentional because providers do not want to make information about topology and resource usage available to

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

By way of example, it is assumed in figure 1, that LSP1 is set up first. Thus
5 CE1 has sent a request to PE1 for a connection to CE2. PE1 uses its local routing database to calculate a label switched path and then sends out signaling messages, using for instance RSVP-TE, to create the calculated path, i.e. LSP1 in the example.

10 It is further assumed that CE1 in a second step intends to set up a second LSP, i.e. LSP2, which shall be disjoint to exclude single points of failure for both LSPs. CE1 will thus send a second connection request to PE2.

However, since CE1 does not participate in the routing protocol on the NNI,
15 it cannot have knowledge about the resources used by LSP1. PE2, on the other hand, has no knowledge about the fact that CE1 is dual-homed via PE1. Therefore, if CE1 just requests a link-disjoint connection from PE2, PE2, cannot know about the existence LSP1 at PE1 and that LSP2, to be established, needs to be resource-disjoint from LSP1.

20 One option would be, that CE1 indicates in the connection request to PE2 that the new connection shall be disjoint from an existing connection established from PE1. PE2 could then directly communicate with PE1 to learn which resources are used by LSP1 and which resources therefore need
25 to be excluded from path calculation for LSP2. Such a direct communication is mentioned in EP1443716B1. However, this would require the definition and implementation of a new protocol for the PE-PE communication, which currently does not exist.

30 The inventors have thus thought of improved mechanisms for supporting LSP diversity within the overlay extension service model considering dual-homing, which build upon existing protocols and allow the CE to provide to the PEs data indicative of shared constraint information.

More particularly, the preferred embodiments use Shared Risk Link Group (SRLG) information as defined in IETF RFC4202, which is incorporated by reference herein.

- 5 A set of links may constitute a Shared Risk Link Group (SRLG) if they share a resource whose failure may affect all links in the set. For example, two fibers in the same conduit would be in the same SRLG. A link may belong to multiple SRLGs. Thus the SRLG Information describes a list of SRLGs that the link belongs to. An SRLG is identified by a 32 bit identifier that is
- 10 unique within an Interior Gateway Protocol (IGP) domain. The SRLG Information is an unordered list of SRLGs that the link belongs to. The SRLG of a LSP is the union of the SRLGs of the links in the LSP. The SRLG of a bundled link is the union of the SRLGs of all the component links.
- 15 Link SRLGs are advertised in the routing protocol and are thus available in the routing database of each PE node. A PE node which processes a path request and performs path calculation can hence determine from its local information the SRLG of the calculated LSP. SRLG information is normally not shared between the transport network and the customer network; i.e.,
- 20 not shared with the CEs of a L1VPN in the L1VPN context.

In the following, two different signaling variations will be described that can be used to allow GMPLS to provide path diversity in a GMPLS controlled transport network.

- 25 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
- 30 with respect to the already established first connection entering the network at a different PE device. The methods according to the preferred embodiments 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.

While the described 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 (i.e. 5 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 10 approaches taken (i.e., in terms of how the provider chooses to perform routing internal to their network).

In a first embodiment, it will be assumed that provider Shared Resource Link Group (SRLG) Identifier information is both available and shareable 15 with the CE. This will be a policy decision of the provider. In the first embodiment SRLG IDs will hence be used as signalled information on a UNI message and passed transparently between PEs via the dual-homed CE

The a second embodiment, it is assumed that provider SRLG IDs are either 20 not available or not shareable (based on provider network operator policy) with the CE. For this case, a mechanism is described where information signalled to the PE on UNI messages does not require shared knowledge of provider SRLG IDs to support LSP diversity for the overlay extension model. Both approaches follow the L1VPN framework defined in RFC4847, which is 25 incorporated by reference herein. While both methods could be implemented in the same PE network, it is likely that an L1VPN CE network would use only one mechanism at a time.

PE SRLG information can be used directly by a CE if the CE understands the 30 context, and the CE view is limited to its L1VPN context. In this case, there is a dependency on the provider information and it is thus preferable 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 customer network. This is possible with both methods detailed below. Even for the method where provider SRLG
5 information is passing through the CE device the two SRLG identifier spaces can remain fully decoupled and the operator of the customer network is free to assign SRLG identifiers from the customer SRLG identifier space to the CE to CE connection that is passing through the provider network. It is worthwhile mentioning that the CE device does not need to process and
10 decode SRLG information.

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
15 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, for example Link state database or Network Management,
20 which is beyond the scope of this document, but which could be implemented by those skilled in the art based on the below description without undue experimentation.

The first embodiment employing an exchange of SRLG information between
25 the PEs via the CE device will now be described in more detail.

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.

30

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. According to the embodiment, CE1 would request the SRLG information of LSP1 from PE1, and then transparently pass this information to PE2 as part of the LSP2 setup request. PE2 is therefore capable of calculating a path
5 for LSP2 that is SRLG disjoint with respect to LSP1.

The exchange of SRLG information is achieved on a per L1VPN 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
10 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 defined in the IETF contribution "RSVP-TE Extensions for Collecting SRLG
15 Information" by F. Zhang et al., draft-zhang-ccamp-srlg-fa-configuration-05.txt, March 2012, which is incorporated by reference herein. However, even if SRLG information is known for several LSPs, only a very limited view of the provider network topology is available from this information.

20 In the first embodiment, the following rules will apply to the CE-PE behavior:

When a dual-homed UNI-C intends to establish an LSP to the same destination UNI-C via another PE node, it can request the SRLG information
25 for an already established LSP by setting the SRLG information flag in the LSP attributes sub-object of the RSVP PATH message.

As long as the SRLG information flag is set in the PATH message, the PE node inserts the SRLG sub-object as defined in draft-zhang-ccamp-srlg-fa-
30 configuration-05.txt into the RSVP RESV message that contains the current SRLG information for the LSP. The PE passes on the SRLG information for the LSP.

If the provider network's policy has been configured so as not to share SRLG information with the customer network, the SRLG sub-object is not inserted in the PATH message even if the SRLG information flag is set.

- 5 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 UNI-C sets the 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
10 receives this request it calculates a path to the given destination and uses the received SRLG information as path computation constraints. If no disjoint LSP can be found, the PE shall return an error message.

- 15 The SRLG diversity flag is new to the RSVP protocol and needs to defined accordingly.

Turning now to figure 2, the communication in the specific example of figure 1 will be the following:

- 20 CE1 has sent an RSVP PATH message (not shown) to PE1 for LSP1, which has the SRLG information flag set. PE1 replies with a RESV Message RM1, which contains the SRLG sub-object of LSP1.

- 25 In a next step, CE1 sends a PATH message PM2 to PE2 to request setup of disjoint path LSP2. CE1 sets the SRLG diversity flag in the LSP attributes sub-object of PATH message PM2 and adds the received SRLG sub-object. PE2 then calculates LSP2 using the received SRLG as exclusion list.

- 30 When PE2 has established LSP2, it returns to CE1 a RESV message RM2, which includes an SRLG sub-object of LSP2. CE1 passes this information on to PE1 in a PATH message PM1.

PE1 and PE2 would then know about the SRLG lists of both LSPs LSP1, LSP2, so that if one of the LSPs has to be redirected, e.g. in a restoration

scenario in case of a failure, the SRLG lists of the other LSP can be used as exclusion list for the path calculation of the redirected LSP.

5 While the first embodiment directly passes the SRLG list for path LSP1 over the UNI to the CE, the second embodiment described below introduces a new identifier called Path Affinity Set identifier, which references the SRLG list and is communicated instead of the latter.

10 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. Optionally, more detailed PATH information of an exclude path or set of paths can be specified. The motive behind the PAS information is to have as little exchange of diversity information as possible between the
15 L1VPN 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 an L1VPN context and therefore only unique with respect
20 to a particular L1VPN.

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

For example, PAS can be used with no prior exchange of PAS information between the CE and the PE. Upon reception of the PAS information the PE can infer the CE's requirements. The actual PAS identifier used will be
30 returned in the RESV message. Optionally an empty PAS identifier allows the PE to pick the PAS identifier. Similar to the first embodiment using 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 that identifier and by association any PE related SRLG information, for any LSPs associated with the resources assigned to the L1VPN. 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 associated with the "123" identifier or optionally another PAS identifier for that L1VPN.

The 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.

The L1VPN Port Information table (PIT) defined in 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 L1VPN basis and contains a PAS identifier, PE addresses and SRLG information.

The CE Path Affinity Set may be used to signal paths without CE Source and Destination addresses; however, the PE will preferably associate the CE SRLG Group with a list of PE SRLG plus the PE addresses associated with this LSP.

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.

Figure 3 shows an Diverse LSP information object that can be used for communication on the UNI between CE and PE. It contains the following fields:

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. It can be used to indicate that there are more than 1 other LSPs, which need to be routed link-diverse.
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 to find a diverse path.

- 5 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 setup or cannot be determined by the PE edge processing the UNI, then the path is only with the Path Affinity set constraint. 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 a mechanism to translate CE Addresses into provider addresses when correlating with provider SRLG information.
- 10
- 15 Other control bits for other semantics can be added as needed. For example if a pure diverse path is not available, find a minimally intersecting path. Or do not allow the PAS identifier to be overridden. The latter could be useful in a case where a CE is configured.
- 20 In the second embodiment, the following rules will apply to the CE-PE behavior:

- When a UNI-C 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 could belong to the same Path Affinity Set. The Path Affinity Set identifier is a value (1 through $2^{32}-255$) that is independent of the mechanism the CE or the PE use for diversity. The Path Affinity Set contains a single identifier that can be used to request diversity and associate diversity. A value of zero is a special value that means the PE may assign a value.
- 25
- 30

When processing a CE PATH message in a L1VPN 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.

- 5 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. If on the other hand, PAS diversity bit is set, the PE would return an error code.

10

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 choose 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.

15

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 (predefined default procedure). An another specific error code would be used to indicate that the PAS object is not understood.

20

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

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

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

The Path Affinity Set information will be distributed in the provider network using a routing protocol. Information about SRLG is already available in the IGP Traffic Engineering database. A PE network can be designed to have additional opaque records for Provider paths that
5 distribute PE paths and SRLG on a L1VPN basis. When a PE path is setup, the following information allows a PE to lookup the PE diversity information:

- L1 VPN Identifier 8 bytes
- Path Affinity Set Identifier
- 10 – Source PE Address
- Destination PE Address
- List of PE SRLG (variable)

The source PE address and destination PE address are the same addresses in
15 the L1VPN 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 L1VPN Identifier is associated with a CE. The only value that is signaled from the CE is the Path Affinity Set and optionally the
20 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. . It may be shared with other PEs or PCE devices through a number of mechanisms.

25 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
30 Affinity Set remains unassigned 0x0000.

Returning now to figure 4, the communication according to the second embodiment would be as follows. First, CE1 will send to PE1 a PATH message PM1' which includes a diversity flag set to 1. Optionally, PATH

message PM1' can also include a proposed value for the PAS identifier. PE1 will calculate and establish LSP1. Path calculation and setup will use well known mechanisms and GMPLS protocols. It should be understood that PATH message PM1' can also be sent in case LSP1 is already active, since
5 according to RSVP, PATH and RESV messages are periodically resent to refreshed the status of an LSP.

When receiving PATH message PM1' with activated diversity flag, PE1 will determine the SRLG of LSP1 and assign a PAS identifier for this SRLG. If
10 PATH message PM1' already contained a value for the PAS identifier, PE1 will preferably use this value. It may be however, that the value is not the best to be used, e.g. because the PE uses a better value. In such case, PE1 will choose a new value and overwrite the value provided by CE1. In any case. PE1 will return to CE1 a RESV message RM1' with a PAS object that
15 contains the chosen value for the PAS identifier.

PE1 will then advertise the previously determined SRLG list and corresponding PAS object via the NNI in the provider network using the OSPF-TE routing protocol. More particularly, PE1 will sent at all its NNIs
20 OSPF messages RTM1 to share the SLRG/PAS information. Each node P1-P4, PE1-PE4 in the provider network will receive the OSPF message RTM1, directly or indirectly (i.e. forwarded), and store the PAS object and corresponding SRLG list in its routing database.

25 In a next step, CE1 will send a path request PM2' to PE2 requesting a new path LSP2. This path request PM2' is a RSVP PATH message which contains the PAS object. The PAS object contains the identifier as received from PE1 as well as the diversity flag D set to 1.

30 PE2 will look up in its routing database the SRLG list that corresponds to the received PAS identifier and use this SRLG list as an exclusion list for the path calculation. The path calculation can thus be performed with only local knowledge of PE2. PE2 will hence calculate LSP2 and initiate path setup using standard RSVP signaling.

As a next step, since the diversity flag was enabled, PE2 will determine the SRLG of LSP2 and assign the same PAS object thereto. PE2 will then send out OSPF routing messages RTM2 to advertise the PAS object with the updated SLRG information. PE1 will receive a routing message RTM2 and update its routing database. Therefore, in case of a failure along LSP1, PE1 is able to restore LSP1 using the SRLG of LSP2 as exclusion list.

PE2 will also return a RESV message RM2' to CE1 to signal the successful creation of LSP2. RESV message RM2' will also contain the PAS object to confirm the PAS identifier, which will be the same as the PAS identifier of LSP1.

A simplified block diagram of a provider edge node PE is shown in figure 5. A provider edge node can be for example a TDM switch, an Optical Cross-Connect (OXC), or a Photonic Cross-Connect (PXC). Alternatively, a PE device may be an Ethernet Private Line (EPL) type of device that maps Ethernet frames onto layer 1 connections by means of Ethernet over TDM.

A provider edge node PE, as shown in figure 5, has one or more User-Network Interfaces UNI, one or more Network-Network Interfaces NNI and a switch device interconnecting user traffic from the UNIs to network connections at the NNIs. The switch device can for instance be for example a TDM switch, a space switch, a wavelength switch or a cell based switch.

The provider edge node PE further has an associated controller CTR, which processes control traffic from the UNIs and the NNIs, and controls the switch device to perform appropriate switch functions. Controller CTR has a processor CPU and a storage device which contains a routing database DB.

Controller CTR runs a GMPLS routing agent and corresponding GMPLS protocol stacks and communicates with attached customer edge node(s) via the UNI(s) and with other provider nodes via the NNI(s). The controller is

programmed to receive and transmit the signaling and routing protocol messages as specified above.

A simplified block diagram of a Customer Edge device CE is shown in figure 6. A CE device can be a variety of devices, for example, Time Division Multiplexing (TDM) switch, router, and layer 2 switch. The Customer Edge device CE has at least two User-Network interfaces to support dual-homing. In the shown embodiment, Customer Edge device CE has one or more Ethernet interfaces ETH for data connection within a customer network and a packet switch PSW selectively interconnecting user packet data from the Ethernet interfaces ETH to traffic paths terminated at the UNIs. The packet switch can be implemented for example using a shared memory switch and a network processor for header processing that steers reading and writing to the memory switch in dependence of destination addresses read from the packet headers.

Similar to provider edge node PE in figure 5, customer edge node CE has an associated controller CTR', which processes control traffic from the UNIs, and controls the packet switch SW to perform appropriate switch functions. Controller CTR' has a processor CPU' and a storage device which stores a routing table DB'. It should be understood that routing table DB' represents the client network routing data base that is distinct from the routing data base in the provider network, of which the client network has no visibility.

Controller CTR runs a UNI agent and corresponding RSVP protocol stack and communicates with attached provider edge nodes via the UNIs. The controller is programmed to receive and transmit the signaling protocol messages as specified above.

The above embodiments make use of a distributed control plane and local GMPLS routing agents running on each provider (edge) node. It should be understood that similar mechanisms and protocol extensions can also be used in a centralized architecture with a Path Computing Element (PCE) and the provider edge nodes PE1, PE2 as Path Computing Clients (PCC).

In general, different PCE/PCC approaches have been proposed and are currently under discussion at IETF. In one approach, the PCE can be a stateful PCE, i.e. having knowledge of all LSPs in the network. In another approach, however, the PCE can be stateless. If a stateless PCE is used, PE1 and PE2 would independently request path computation from the PCE include the diversity attributes in their path computation requests. The PCE will then perform path computation for the PEs, one after the other, will additionally determine and return SRLG information and consider the SRLG information as exclusion list to achieve link-diverse LSPs.

More particularly, in the first embodiment, when a PCE is involved, PE1 will receive the SRLG of LSP1 from the PCE and forward it in return message RM1 to CE1. PE2, upon receipt of path request message PM2, which includes the SRLG list that corresponds to LSP1, will forward the SRLG information to the PCE so that the PCE can this as exclusion list for the calculation of LSP2. This would enable even a stateless PCE to perform link-diverse path calculation in a dual-homing network topology. In a similar way, also the second embodiment can be employed in a context with a stateless PCE.

A person of skill in the art would readily recognize that steps of various above-described methods can be performed by programmed computers. Herein, some embodiments are also intended to cover program storage devices, e.g., digital data storage media, which are machine or computer readable and encode machine-executable or computer-executable programs of instructions, wherein said instructions perform some or all of the steps of said above-described methods. The program storage devices may be, e.g., digital memories, magnetic storage media such as a magnetic disks and magnetic tapes, hard drives, or optically readable digital data storage media. The embodiments are also intended to cover computers programmed to perform said steps of the above-described methods.

The description and drawings merely illustrate the principles of the invention. It will thus be appreciated that those skilled in the art will be able to devise various arrangements that, although not explicitly described or shown herein, embody the principles of the invention and are included within its spirit and scope. Furthermore, all examples recited herein are principally intended expressly to be only for pedagogical purposes to aid the reader in understanding the principles of the invention and the concepts contributed by the inventor(s) to furthering the art, and are to be construed as being without limitation to such specifically recited examples and conditions. Moreover, all statements herein reciting principles, aspects, and embodiments of the invention, as well as specific examples thereof, are intended to encompass equivalents thereof.

The functions of the various elements shown in the figures, including any functional blocks labeled as “processors” or “CPUs”, may be provided through the use of dedicated hardware as well as hardware capable of executing software in association with appropriate software. When provided by a processor, the functions may be provided by a single dedicated processor, by a single shared processor, or by a plurality of individual processors, some of which may be shared. Moreover, explicit use of the term “processor” or “controller” should not be construed to refer exclusively to hardware capable of executing software, and may implicitly include, without limitation, digital signal processor (DSP) hardware, network processor, application specific integrated circuit (ASIC), field programmable gate array (FPGA), read only memory (ROM) for storing software, random access memory (RAM), and non volatile storage. Other hardware, conventional and/or custom, may also be included. Similarly, any switches shown in the FIGS. are conceptual only. Their function may be carried out through the operation of program logic, through dedicated logic, or through the interaction of program control and dedicated logic, the particular technique being selectable by the implementer as more specifically understood from the context.

Claims

5

1. A method of establishing traffic paths through a provider network, wherein a customer edge device (CE1) is connected to a first provider edge node (PE1) and to a second provider edge node (PE2) via User-Network Interfaces (UNI); wherein said provider edge nodes (PE1, PE2) are
- 10 interconnected directly or indirectly with provider core nodes (P1-P4) and further provider edge nodes (PE3, PE4) via Network-Network Interfaces (NNI) to form said provider network; and wherein said provider edge nodes (PE1-PE4) and said provider core nodes (P1-P4) exchange routing information over said Network-Network Interfaces (NNI) using a routing
- 15 protocol; the method comprising:
- requesting by said customer edge device (CE1) a first traffic path through said provider network to be established from said first provider edge node (PE1);
 - within said provider network, calculating said first traffic path (LSP1)

20 using said routing information and initiating setup of said first traffic path (LSP1) using a signaling protocol;

 - requesting by said customer edge device (CE1) a second traffic path (LSP2) through said provider network to be established from said second provider edge node (PE2) and to be disjoint from said first traffic path

25 (LSP1); and

 - within said provider network, calculating said second traffic path (LSP2) using said routing information and a constraint to be disjoint from said first traffic path (LSP1) and initiating setup of said second traffic path (LSP2) using said signaling protocol;
- 30 further comprising:
- determining shared constraint information of said first traffic path (LSP1) and returning from said first provider edge node (PE1) to said customer edge device (CE1) via said User-Network Interface a data element indicative of said shared constraint information; and

- forwarding said data element from said customer edge device (CE1) to said second provider edge node (PE2) to enable path calculation of said second traffic path using said shared constraint information as exclusion list;
- 5 wherein said shared constraint information comprises Shared Risk Link Group identifiers of all provider network links used by said first traffic path (LSP1); and wherein said data element comprises an unordered list of Shared Risk Link Group identifiers of said first traffic path.
- 10 2. A method according to claim 1, wherein said step of determining said shared constraint information is carried out upon receipt of a message sent by said customer edge device (CE1) to said first provider edge node (PE1) having a diversity flag set.
- 15 3. A method according to claim 2, wherein said message is a PATH message according to the Resource ReserVation Protocol with Traffic Engineering Extension.
- 20 4. A method according to claim 1, wherein said step of returning said data element to said first customer edge device (CE1) comprises sending a RESV message according to the Resource ReserVation Protocol with Traffic Engineering Extension comprising said data element as an LSP attributes sub-object.
- 25 5. A method according to claim 1, comprising
- after calculating said second traffic path (LSP2) determining at said second provider edge node (PE2) shared constraint information of said second traffic path (LSP2) and returning to said customer edge device (CE1) via said User-Network Interface a second data element indicative
- 30 of said shared constraint information of said second traffic path (LSP2); and
- forwarding said second data element from said customer edge device (CE1) to said first provider edge node (PE1).

6. A method according to claim 1, wherein said network further comprises one or more Path Computing Elements which communicate with said provider edge nodes (PE1-PE4) and wherein said steps of calculating traffic paths and determining shared constraint information for said traffic paths
5 are carried out by any of said Path Computing Elements.

7. A provider edge node comprising at least one User-Network Interface (UNI) for connecting to a customer edge device (CE1), at least one Network-Network Interface (NNI) for connecting to other provider nodes, a
10 switch device (SW) interconnecting user traffic from the User-Network Interface (UNI) to network connections at the Network-Network Interface (NNI), and a controller (CTR); said controller (CTR) comprising a programmable processor (CPU) and a routing database (DB); wherein said controller is programmed to
15 – exchange routing information over said Network-Network Interfaces (NNI) using a routing protocol and storing routing information received at said Network-Network Interfaces (NNI) in said routing database (DB)
– receive via said User-Network Interface (UNI) a request to establish a traffic path through said provider network;
20 – calculate said traffic path (LSP1) using said routing information and
– initiating setup of said traffic path (LSP1) using a signaling protocol; and
– determine shared constraint information of said traffic path (LSP1) and return to said customer edge device (CE1) via said User-Network Interface a data element indicative of said shared constraint
25 information;

wherein said shared constraint information comprises Shared Risk Link Group identifiers of all provider network links used by said first traffic path (LSP1); and wherein said data element comprises an unordered list of Shared Risk Link Group identifiers of said first traffic path.



30

8. A provider edge node according to claim 13 wherein said controller is further programmed to receive together with a request to establish a traffic path a data element indicative of said shared constraint information

and to perform path calculation using said shared constraint information as exclusion list.

9. A customer edge device, comprising at least two User-Network
5 interfaces (UNI) for connecting to two different provider edge nodes (PE1, PE2) and a controller (CTR'), said controller is programmed to
- request via a first of said User-Network interfaces (UNI) a first traffic path (LSP1) through said provider network;
 - receive at said first User-Network interface (UNI) a data element
10 indicative of shared constraint information of said first traffic path (LSP1); and
 - request via a second of said User-Network interfaces (UNI) a second traffic path (LSP2) through said provider network to be disjoint from said first traffic path (LSP1); and
 - 15 – forward at said second User-Network Interface (UNI) said data element received via said first User-Network Interface (UNI) to enable path calculation of said second traffic path (LSP) using said shared constraint information as exclusion list;
- wherein said shared constraint information comprises Shared Risk Link
20 Group identifiers of all provider network links used by said first traffic path (LSP1); and wherein said data element comprises an unordered list of Shared Risk Link Group identifiers of said first traffic path.

1/3

**Fig. 1****Fig. 2**

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																	0	C-Type															
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																			
ADDR Length									No. of PAS									D	reserved																															
Path Affinity Set identifier																																																		
Source Address (variable)																																																		
Destination Address (variable)																																																		

Fig. 3






Fig. 4

3/3

PE**Fig. 5**CE**Fig. 6**

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2013/062501

A. CLASSIFICATION OF SUBJECT MATTER
INV. H04L12/735
ADD.

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

EPO-Internal, INSPEC, WPI Data

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 2009/103442 A1 (DOUVILLE RICHARD [FR]) 23 April 2009 (2009-04-23) paragraphs [0003] - [0022], [0046] - [0056]	1-9
Y	----- EP 2 267 954 A1 (ALCATEL LUCENT [FR]) 29 December 2010 (2010-12-29) cited in the application paragraphs [0004], [0012], [0019] - [0026], [0038] - [0044]	1-9
Y	----- EP 1 443 716 B1 (CIT ALCATEL [FR] ALCATEL LUCENT [FR]) 2 May 2007 (2007-05-02) cited in the application paragraphs [0009] - [0011], [0023] - [0031] ----- -/-	1-9



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

18 September 2013

Date of mailing of the international search report

25/09/2013

Name and mailing address of the ISA/

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040,
Fax: (+31-70) 340-3016

Authorized officer

Ceccarini, Giulio

INTERNATIONAL SEARCH REPORT

International application No
PCT/EP2013/062501

C(Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	"Architecture and specification of data communication network; G.7712/Y.1703 (09/10)", ITU-T STANDARD, INTERNATIONAL TELECOMMUNICATION UNION, GENEVA ; CH, no. G.7712/Y.1703 (09/10), 6 September 2010 (2010-09-06), pages 1-96, XP017467837, page 80 - page 81 -----	1-9
X,P	FEDYK D BELLER LIEVEN LEVRAU ALCATEL-LUCENT D: "Layer 1 VPN Enhanced Mode - Overlay Extension Service Model; draft-fedyk-ccamp-l1vpn-extnd-overlay-00.txt", LAYER 1 VPN ENHANCED MODE - OVERLAY EXTENSION SERVICE MODEL; DRAFT-FEDYK-CCAMP-L1VPN-EXTND-OVERLAY-00.TXT, INTERNET ENGINEERING TASK FORCE, IETF; STANDARDWORKINGDRAFT, INTERNET SOCIETY (ISOC) 4, RUE DES FALAISES CH- 1205 GENEVA, SWITZERLAND, 9 July 2012 (2012-07-09), pages 1-19, XP015083856, the whole document -----	1-9

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No

PCT/EP2013/062501

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2009103442	A1	23-04-2009	AT 480075 T 15-09-2010
		CN 101399771 A 01-04-2009	
		EP 2043310 A1 01-04-2009	
		ES 2351870 T3 11-02-2011	
		FR 2921779 A1 03-04-2009	
		US 2009103442 A1 23-04-2009	
		WO 2009040421 A1 02-04-2009	

EP 2267954	A1	29-12-2010	AT 527790 T 15-10-2011
		CN 102804710 A 28-11-2012	
		EP 2267954 A1 29-12-2010	
		JP 2012531155 A 06-12-2012	
		KR 20120019485 A 06-03-2012	
		US 2012089673 A1 12-04-2012	
		WO 2010149562 A1 29-12-2010	

EP 1443716	B1	02-05-2007	AT 361611 T 15-05-2007
		DE 60313572 T2 27-12-2007	
		EP 1443716 A1 04-08-2004	
		US 2005021873 A1 27-01-2005	
