CCAMP Working Group — Dimitri Papadimitriou (Alcatel)
Category: Internet Draft — Fabrice Poppe (Alcatel)
Expires: May 2003 — Sudheer Dharanikota (Consult)
Riad Hartani (Caspian)
Raj Jain (Nayna)
Jim Jones (Alcatel)
Senthil Venkatachalam (OPNet)
Yong Xue (WorldCom)

November 2002

Shared Risk Link Groups Encoding and Processing

draft-papadimitriou-ccamp-srlg-processing-01.txt (*)

Status of this Memo

(*) Formerly draft-many-inference-srlg-02.txt

Abstract

The Shared Risk Link Group (SRLG) concept introduced in [IPO-FRM] and extended in [CCAMP-SRG] is considered as one of the most important criteria concerning the constraint-based path computation of optical channel routes. By applying the SRLG disjointness criteria to the constraint-based path computation, one can select a route taking into account both resource and logical structure disjointness. That implies a lower probability of simultaneous optical channel failure.

This document describes the various physical and logical resource types considered in the SRLG concept. The proposed method focuses on the inference of SRLG information between the network physical and logical layers. The main applications of the proposed model are related to the Constraint-based Shortest Path First (CSPF) algorithm for optical channel route computation and the aggregation of the SRLG information flooded throughout traffic engineering extensions of the IGP routing protocols (such as OSPF and IS-IS).

1. Introduction

Many proposals include the Shared Risk Link Group (SRLG) concept when considering the disjointness for the constraint-based path computation of optical channel routes. In the optical domain, the SRLG concept is used for deriving a path, which is disjoint with respect to physical (or passive) and logical (or active) resources. [CROCH] describes algorithms that can use the SRLG information inferred using the mechanisms described in this document, in order to determine survivable logical topologies on top of a given physical topology.

The corresponding requirements have already been described in [IPO-IMP] while considering physical network topology and associated risks. In the scope of this document, these requirements can be summarized as follows:

1. The SRLG encoding mechanism should reduce the path computation complexity.
2. The SRLG information flooding should be scoped to reduce the amount of information that is exchanged across domains.
3. The SRLG encoding should accommodate the physical and logical restrictions imposed on the diversity requirements.

However, the definition of SRLG in the current format described in [GMPLS-OSPF] and [GMPLS-ISIS] does not provide:

1. The relationship between physical (or logical) resources and between physical and logical resources. For example, a fiber could be part of a sequence of fiber segments or an optical channel may cover several fibers links.
2. The risk assessment during path computation implying the assignment of a conditional failure probability to the SRLGs.
3. The analysis of the specifics aspects of constraint-based path computation and path re-optimization taking SRLG information into account.

This document proposes among others a technique to compute the SRLG with respect to a given risk type. This is achieved by identifying for a given physical layer the resources belonging to an SRLG. The proposed model also permits to compute the dependencies of these resources to the lower (passive) physical layers. The result of the computation also enables to determine the risk associated to each of the SRLGs, that is, the probability that two entities (both using network resources belonging to a given SRLG) go simultaneously down if one of them goes down (thus defining a conditional failure probability).

The remainder of this memo is organized as follows. In section 3, we present the hierarchical model of the network resources and the corresponding SRLG encoding. In section 4, we discuss the use of the proposed approach for risk assessment during path computation and selection. Considerations on topology summarization and path disjointness with respect to SRLG are proposed in section 5.

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 [1].

The reader is assumed to be familiar to the terminology and concepts detailed in [GMPLS-RTG], [GMPLS-OSPF] and [GMPLS-ISIS].

3. Hierarchical Model

The proposed approach is based on a representation of the network resources in terms of a hierarchy. This hierarchical structure is related to the fiber topology (or more generally the physical resources) of the optical network and the channels built on top of this physical topology.

The encoding of the SRLG could be either mapped on this hierarchical model or simply use a flat encoding scheme. Difference between both encoding relies on the extended usage of the SRLG identifiers in the context of diverse route computation (i.e. path disjointness). Since a link can belong to more than one SRLG, an SRLG identifier list (in the SRLG Sub-TLV), as described in [GMPLS-OSPF] and [GMPLS-ISIS] is associated with each link (i.e. the SRLG Sub-TLV is defined as a Sub-TLV of the TE Link TLV). This results in a linear, unordered and non-structured information from which the underlying structure cannot be deduced.

Consequently, either a type field indicating the type of network resource (or logical structure) to which this SRLG identifier refers extends the flat encoding scheme or the encoding itself translates the underlying hierarchical structure. Worth mentioning here is that a hierarchical encoding (since depending on the physical layer which is static by definition) needs an additional mapping structure in order to keep the relationship with link identifiers.

3.1 Network Resource Hierarchy

The network (physical and logical) resource model considered in the inference of the Shared Risk Link Groups (SRLGs) is based on premises detailed in [IPO-FRM] and [IPO-IMP]. The network resource hierarchy includes the logical resources built on top of the physical resource hierarchy belonging to the optical network.

The concepts around network resource hierarchy are based on the following logical resource definitions:

- Sub-Channel: a dedicated container included within a given optical channel uniquely identifies a sub-channel.

- Channel(*): an optical channel is uniquely identified by its corresponding and dedicated wavelength (sub-)interface identifier.

(*) when conversion or amplification is non-collocated with the corresponding interface capability, the passive components are logically embedded into the fiber link (see below definition).

These resources are built on top of the physical resource hierarchy composed by the following physical resources considered as a common denominator of most Optical Transport Network (OTN) environments:

- Fiber Link: a fiber connects two node ports communicating through one or more optical channel if their interfaces support (Dense) Wavelength Division Multiplexing û (D)WDM, also a fiber link can be decomposed as a sequence of fiber spans separated by optical amplifiers

- Fiber Sub-segment: a group of fiber links, any fiber link of a sub-segment start and terminates at the same location

- Fiber Segment: a group of fiber sub-segments, any fiber sub-segment of segment start and terminates at the same location

Moreover, one can consider sequences of fiber (sub-)segment starting and terminating at the same node and defined them as fiber trunks (also referred to as cable or fiber duct).

Thus, the approach developed here extends the definition and concepts given in [IPO-FRM] and [IPO-IMP] by enabling fiber topologies not strictly limited to point-to-point physical inter-connections between nodes.

As represented in Figure 1, the fiber trunk going from location N[1] to location N[3] is composed by the fiber segments A and B and the fiber trunk from location N[1] to location N[2] includes the fiber segments A, C and D.

```
  Location N[1]                                       Location N[3]

Sub-Segment A[1]                                    Sub-Segment B[1]
  ---------------------------------------------------------------
=== . . . ====== Fiber                     Fiber ====== . . . ====
=== . . . ====== Fiber                     Fiber ====== . . . ====
  --------------------------------------------------------------
    Segment A                                          Segment B
  ------------------------------   -----------------------------
=== . . . ====== Fiber          | |        Fiber ====== . . . ====
=== . . . ====== Fiber          | |        Fiber ====== . . . ====
  -------------------------     | |     -------------------------
```

```
   Sub-Segment A[n]           |    | |    |          Sub-Segment B[n]
                              |    | |    |
                              |    | |    | Segment C
                              |    | |    |
   Sub-Segment D[1]           |    | |    |          Sub-Segment E[1]
     -------------------------     | |     -------------------------
   === . . . ====== Fiber          | |        Fiber ====== . . . ====
   === . . . ====== Fiber          | |        Fiber ====== . . . ====
     ------------------------------   ------------------------------
       Segment D                                         Segment E
     ---------------------------------------------------------------
   === . . . ====== Fiber                     Fiber ====== . . . ====
   === . . . ====== Fiber                     Fiber ====== . . . ====
     ---------------------------------------------------------------
   Sub-Segment D[n]                                  Sub-Segment E[n]

     Location N[2]                                     Location N[4]
```

Figure 1. Example of Physical Topology

In this figure, the Segment A is composed by the fiber sub-segments A[1], A[2], ..., A[i], ..., A[n]. The same terminology applies for the segments B, C, D and E.

Consequently, the fiber trunk from location N[2] to location N[4] includes the sub-segments D[2] to D[n] and their corresponding sub-segments within the segment E: E[2] to E[n]. The fiber trunk from location N[1] to location N[2] includes the fiber sub-segments A[n], C[1] and D[1].

3.3 SRLG Definition and Properties

A SRLG is defined as the set of links (or optical lines) sharing a common physical resource (including fiber links, sub-segment and segment) i.e. sharing a common risk. For instance, a set of links L belongs to the same SRLG s, if established over the same fiber link l.

3.3.1 SRLG Properties

The SRLG properties can be summarized as follows:

1) A link belongs to more than one SRLG if and only if it crosses (at least) one of the resources covered by each of these SRLGs

For instance: fiber link l belongs to the SRLG s1 and s2, if and only if it crosses the fiber sub-segment A[1] and B[1]

2) Two links belonging to the same SRLG can belong individually to other (one or more) SRLGs

For instance: fiber links l1 and l2 belong to SRLG s3 (segment A) while l1 belongs to SRLG s1 (since it covers sub-segment A[1]) and l2 to SRLG s4 (since it covers sub-segment D[1])

3.4.2 SRLG Disjointness

To define the LSP SRLG disjointness notion, we first introduce the following definition: an LSP (i.e. sequence of links) covers an SRLG if and only if it crosses one of the links belonging to that SRLG. For instance: LSP p1 covers SRLG s1 (since it crosses link l1)

LSP SRLG disjointness can then be defined as follows.

1) Two LSPs are disjoint with respect to an SRLG s1 if and only if at most one of them covers this SRLG.

For instance: LSPs p1 and p2 are disjoint with respect to SRLG s1 since only p1 covers SRLG s1.

2) Two LSPs are disjoint with respect to a set of SRLG S if and only if the sets of SRLGs they individually cover are completely disjoint

For instance: LSP p1 and p3 are disjoint with respect to set of SRLG S = {s1, s2, s3} since only p1 covers SRLG set S.

3.5 Computational Model Capabilities

This section briefly describes guidelines for an SRLG computational model described in appendix A and based on the above definitions. The main features of this model are:

- Support Constraint-based Shortest Path First (CSPF) algorithm for explicit route (or path) computation by considering SRLG disjointness with respect to one or more risk types

- Encompass hierarchical dependencies between physical resources (inference of SRLG sets using bottom-up computation)

- CSPF computation including the relationship between physical resources and topological structures. For instance:
  . a fiber link can be part of given fiber trunk
  . a fiber cable passing through an earthquake region

- Provide risk assessment during path computation implying allocation of conditional failure probabilities with SRLGs

- SRLG information flooding well-scoped to reduce the amount of link-state advertisements by using summarization

Consequently, the above features suggest an SRLG encoding mechanism that enables:
- Accommodation of the physical and topological hierarchies
- Reduction of the path (i.e. explicit route) optimality

3.6 SRLG Encoding

Using the above definitions and properties, the objective of the hierarchical encoding is to achieve aggregation of the SRLG Identifiers on top of the (passive) optical network topology. For this purpose, we propose a linear encoding scheme including a type field. This provides abstraction of the physical layer structure and should facilitate the management of the SRLG Identifiers.

The proposed SRLG sub-TLV of the TE Link top level TLV includes the following SRLG Identifier (64-bit field):

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      Type     |                    Weight                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Identifier                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Within the SRLG Identifier field (64 bits), the Type field (8 bits) defines the resource type (also generically referred to as the link type) to which the Identifier (32 bits) integer value refers. The weight field (24 bits) assigned on a per-SRLG basis along with the Identifier is defined in Section 4.

The following Type field values are currently defined for physical resources, any other value being reserved:

Type	Value
Reserved	0x00
Fiber Trunk	0x10
Fiber Segment	0x20
Fiber Sub-segment	0x30
Fiber Link	0x40
Node(*)	0xFF

(*) represents a non-GMPLS capable switching elements

Logical resources such as optical channels and TDM circuits (or optical sub-channels) can take one of the following values:

Type	Value
Optical Channel	0x50
Optical Sub-Channel	0x60

Since a given physical or logical resource (for instance a fiber link) can belong to more than one SRLG, the SRLG Identifier structure is defined in the general case as a list of SRLG Identifier (n x 64-bit):

```
 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
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      Type     |                    Weight                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Identifier                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      Type     |                    Weight                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Identifier                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
//                             ...                             //
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|      Type     |                    Weight                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                           Identifier                          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Therefore, though we propose a linear encoding, the aggregation of the SRLG is still possible. This encoding also enables to perform summarization (which is not equivalent to aggregation) at the boundaries of structures defining the spatial coverage of an SRLG Identifier list while overcoming the drawbacks of full hierarchical encoding schemes. In this context, the weight value follows the rules defined in the Section 4.

4. Risk Assessment

Risk assessment is defined as the quantification process of the potential risk associated to the inclusion of a given resource (belonging to a given resource type) located within a given logical structure such as a geographical location for a given logical network resource such as an optical channel.

4.1 Rationale for Risk Assessment

Consider the following example, where the client device makes the following connection requests to the optical network:

- Request for a persistent connection with 99.999 % percent of availability or equivalently a downtime less than X minutes per year.

- Request for protection for a portion of the traffic (at the expense of more charging) compared to other portion of the traffic left unprotected and considered as extra traffic by the server.

Such requirements will be translated into path specific requests. Such path specific requests can be grouped into path selection requirements and path characterization requirements.

1. Path selection requirements

These typically dictate which physical path should be taken to achieve the availability requirements of the client request. These requirements are typically related to the logical and physical diversity as mentioned in Section 3.

2. Path characterization requirements

Path characterization requirements typically dictate the recovery mechanisms as specified by the client connection request. This can be achieved in the form of section and/or path recovery but also depending on the ring or meshed topology of the optical network. However, these considerations are out of the scope of this document.

The components that need formalization in this example are:
- Step 1. Specification of the client requirements (service level).
- Step 2. Configuration of the network that helps in assessing the service specifics such as the availability.
- Step 3. Propagation of the above-configured information.
- Step 4. Processing of the above-propagated information.

Step 1 of specifying the requirements is not in the scope of this document. Steps 2 to 4 are discussed in the following sections where we elaborate on the risk assessment during path selection.

4.2 Risk Assessment

Risk (the complementary of availability) assessment is defined as the evaluation of the potential risk associated to the inclusion for a given path of a specific network resource (i.e. belonging to a given resource type). Given that an SRLG Identifier list is used to encode the group of logical and physical resources, if a mechanism is devised to assign the risk of failure associated with the corresponding resource, the availability of the path using these resources can be computed. This, in order to meet the connection availability as requested by the client request.

A simple approach is to assign the conditional failure probability to each of the SRLG Identifiers. This information is encoded in the weight (24 bits) field along with the SRLG information as defined in Section 3.3. The associated weights can be used to either increase or decrease the potential usage of the resource (i.e. inclusion into the selected route). The Weight field is defined as a 24-bit length encoded integer number (maximum value 2^24 - 1) to be divided by (2^24 û 1) and multiplied by 100 to get the corresponding value in percents. For instance a conditional failure probability of 0.99999 corresponds to a weight of 16.777.047 and 0.00005 to 839.

For instance, we can associate a conditional failure probability of 25% (weight = 4.194.304) to any fiber sub-segment. It means that by selecting two (or more than two) different optical channel paths including the same SRLG identifier with respect to fiber sub-segment failure, if one of these connections fails, then the probability that the other connection fails is 25%.

The failure probability of a fiber can also depend on the length of the fiber. In addition, a fiber can pass across different regions with different failure probabilities. In this case, we need to consider an aggregated failure probability per fiber taking into account each of the failure probabilities of its sub-components.

For instance, if we refer to our previous example and by considering

1. a conditional failure probability of 5% is associated to any fiber link
2. a conditional failure probability of 1% is associated to any fiber segment

Then by selecting two different optical channels included within the same SRLG with respect to fiber segment failure, we obtain a simultaneous connection failure probability of 1%. Consequently, when a protected connection service is requested, by choosing fiber segment path disjointness, the simultaneous connection failure probability is also of 1%. However, if two optical channels flowing through the same fiber link are selected, then we have a probability of 5% that both optical channels fail simultaneously.

More generically, if one assigns a conditional failure probability c[i] to each SRLG Identifier i, the resulting conditional failure probability C for a logical resource (a connection, for instance) using these N resources is given by:

C = 1 - (1 û c[1]) x (1 û c[2]) x . . . x (1 û c[N])

When all the conditional failure probabilities c[i] are small (i.e. c[i] << 1) C can be approximated by (1 - Sum[i]c[i]).

4.3 Risk Assessment Application

The relationship between the availability of the service requested by the client (i.e. a working and a protection connection) and the conditional failure probabilities of the physical resource elements included within the corresponding paths can be described as follows.

If we consider for instance 1) a conditional failure probability of 1% if fiber links are selected within the same fiber trunk 2) a conditional failure probability of 1% if fiber links are selected within the same fiber segment and 3) the corresponding failure are independent events.

Then the availability of a connection whose backup is established over fiber links (in same fiber segment) and within the same fiber trunks as the primary connection itself is approximately 98% (by using the above formula).

Note that the initial value of the conditional failure probability needs to be statically encoded; however, based on the history of the failures these values could be dynamically re-evaluated. The corresponding mechanism needs to be specified and is left for further study.

5. Applications

The SRLG method developed here results in applications related to the CSPF explicit route computation and the SRLG identifier sets summarization. The latter enables in turn, intra- and inter-area diverse path computation. For that purpose we first extend the SRLG concept for logical resources such as optical channels and optical sub-channels (i.e. TDM connections).

5.1 Extension to the Logical Resources

The SRLG concept can be extended to logical structures and resources by taking into account the following:

1. Given the physical decomposition of the optical network topology, the SRLG encoding can be hierarchically structured. The hierarchical encoding helps in constructing the logical topology abstraction, which in turn can be used in SRLG summarization and loose-path computation. The link semantic may also be extended to accommodate virtual links.

2. Propagate the SRLG information related to these logical (structures and resources) links, for instance (a set of) optical channels using IGP flooding for intra- and inter-area routing purposes.

3. Reduce the amount of the flooded information and hence explicit route computation optimality and extend the flooding scope of the propagated information to accommodate logical resources (i.e. optical channels and TDM circuits).

5.2 SRLG Information Flooding and Summarization

The SRLG set of each link (i.e. physical and logical resources) is encoded as described in Section 3.3. This information is propagated once at initial configuration between the various nodes using the GMPLS Traffic Engineering extensions to IGPs (see [GMPLS-OSPF] and IS-IS [GMPLS-ISIS]). After this initial SRLG identifier exchange, corresponding values do not change over time.

Nevertheless, the propagation of SRLG information will be necessary whenever a new link is added or an existing link is removed. Initially, it is assumed that the failure probabilities of the various resources are (statically) configured. However, it is envisioned that after a certain running period, the failure probability associated to the SRLG and propagated along with the SRLG sub-TLV itself (as described in Section 3.3) will change over time and thus give rise to a dynamic metric.

5.3 Bottom-Up Computation of SRLG Relationships

Once the traffic engineering (topology-related) link information is received by the node, the corresponding SRLG relationships can be computed on a regular basis.

The fiber trunk SRLG relationships are used to compute the fiber segment SRLG relationships, which in turn are used to compute the fiber sub-segment SRLG relationships and finally the fiber SRLG relationships. To each SRLG relationship (at each level), which defines the membership of a resource to a particular SRLG, we associate the conditional failure probability between two resources belonging to this level (for instance, between two fibers that have a common fiber sub-segment).

5.4 Topology Summarization and Route Computation

A direct application of this model is the Constraint-based Shortest Path First (CSPF) algorithm used for explicit route computation (i.e. traffic-engineered LSP creation) to maximize the connection disjointness and so decrease their common failure probability. Given an existing set of connections across the network, the objective is to compute, for a newly requested connection, an explicit route across the optical network topology such that this connection is diversely routed from an existing given set of connections.

The diversity requirement is a routing constraint, that can be expressed in terms of a conditional failure probability with respect to an existing (set of) connections and more generally any kind of resources. Hence, in addition to the other traffic-engineering constraints, the diversity constraint requires that the conditional failure probability does not exceed a given threshold. The CSPF algorithm needs to be updated to take the routing diversity constraint into account.

The SRLG concept generates another dimension to the existing constraint-based path computation methods traditionally used in hierarchical networks. The SRLG constraints comes in addition to the common traffic-engineering constraints such as bandwidth availability, link metrics and TE attributes (see [OSPF-TE]). The specificity of the routing diversity constraint requires the use of a path computation algorithm that provides not only complete multi-path disjointness but also partial multi-path disjointness with respect to various risk factors. In a similar way, appropriate mechanisms should also be used in order to perform path re-optimization following various recovery strategies.

The specific aspect of complete and partial multi-path disjointness is related to the fact that a path may with respect to SRLG be fully or partially disjoint from a given set of path. In brief, one speaks about SRLG partial or full multi-path disjointness. This means that a connection may be disjoint from another but only to some extent with respect to some risk factors. Thus when referring to path disjointness with respect to SRLG one may also include the ratio of disjointness with respect to various risk type assigned to their component resources. Consider for instance the LSPs p1 and p2 using j1 and j2 resources respectively, m of these resources having an SRLG in common, with m =< j1 and m =< j2, then [(j1-m) + (j2-m)]/(j1+j2) provides the disjointness ratio. For instance, if j1=13, j2=7 and m=1, this ratio equals 0.90, thus corresponding path p1 and p2 are SRLG disjoint at 90%.

As such, this feature differentiates path disjointness constraint from any other constraint commonly considered in path computation.

6. Scalability Considerations

6.1 OSPF Scalability Considerations

A node SHOULD minimize the amount of routing information it floods. Each time a SRLG sub-TLV is configured or removed that information shall be flooded (not necessarily immediately) to all nodes in the routing domain. This results in updating an existing LSA or flushing an existing LSA. Removing an LSA from the (TE) Link State DataBase (TE LSDB) can be accomplished in OSPF by prematurely aging the LSA. The LSA is re-flooded with an LSA age equal to MaxAge. Each node receiving an existing LSA with MaxAge removes it from its link state database.

Also, the usage of OSPF implies each LSA must be refreshed periodically (when the LSA age field reaches the LSRefreshTime, see [RFC-2328]) to avoid age timeout and removal from the link state database. This periodical LSA flooding and processing applies particularly to the TE link capability sub-TLVs defined in this document since their variation period is expected to be much larger than the LSRefreshTime.

An Opaque LSA has a length field of 16 bits indicating the length of the LSA, including the header. Thus, the length of OSPF packets can be up to 65535 octets (including the IP header). Moreover, an OSPF packet can contain several (Opaque) LSAs. OSPF relies if necessary on the IP fragmentation to transmit large packets without any loss of functionality. However this is not recommended and it is suggested to split packets that are too large into several smaller packets.

Therefore, the proposed SRLG sub-TLVs defined in this document (and included in the top level TE Link TLV) MUST not exceed the maximum OSPF packet size. This limits the number of SRLG identifiers that a sub-TLV can include to approximately 125 (also this number largely depends on the other sub-TLVs included in the corresponding TE LSA constraining to take for the sake of this application a conservative approach).

6.2 IS-IS Scalability Considerations

TBD.

7. Compatibility Considerations

7.1 OSPF Compatibility Considerations

There should be no interoperability issues with GMPLS-capable nodes that do not implement the proposed SRLG extensions since the sub-TLVs proposed in this document is optional and thus will be silently ignored.

The result of having such nodes that do not implement the proposed SRLG extensions is largely detailed in this memo. However, SRLG constraint paths can still be calculated using the SRLG sub-TLVs as proposed in [GMPLS-OSPF].

The present memo mandates these sub-TLVs to be mutually exclusive i.e. they MUST never appear simultaneously as sub-TLV of the same top level Link TLV [OSPF].

7.2 IS-IS Compatibility Considerations

TBD.

8. Security Considerations

Security considerations related to SRLG and related applications are left for further study.

9. References

[CROCH] O.Crochat, J.-Y. Le Boudec and O. Gerstel, "Protection Interoperability for WDM Optical Networks", IEEE/ACM Transactions on Networking, Vol. 8, No. 3, June 2000, pp. 384-395.

[IEEE-ORL] J.Strand et al., æIssues for Routing in the Optical LayerÆ, IEEE Communication Magazine, Volume 39, Number 2, FebÆ01.

[GMPLS-ISIS] K.Kompella et al., æISIS Extensions in Support of Generalized MPLSÆ, Internet Draft, Work in Progress, draft-ietf-isis-gmpls-extensions-14.txt.

[GMPLS-OSPF] K.Kompella et al., æOSPF Extensions in Support of Generalized MPLSÆ, Internet Draft, Work in Progress, draft-ietf-ccamp-ospf-gmpls-extensions-08.txt.

[GMPLS-RTG] K.Kompella et al., æRouting Extensions in Support of Generalized MPLS,Æ Internet Draft, Work in Progress, draft-ietf-ccamp-gmpls-routing-05.txt.

[IPO-FRM] J.Luciani et al., æIP over Optical Networks A FrameworkÆ, Internet Draft, Work in progress, draft-ietf-ipo-framework-03.txt.

[IPO-IMP] J.Strand et al., æImpairments And Other Constraints On Optical Layer RoutingÆ, Internet Draft, Work in progress, draft-ietf-ipo-impairments-02.txt.

[MPLS-BDL] K.Kompella et al., æLink Bundling in MPLS Traffic EngineeringÆ, Internet Draft, Work in progress, draft-ietf-mpls-bundle-04.txt.

[OSPF-TE] D.Katz, D.Yeung and K.Kompella, ôTraffic Engineering Extensions to OSPFö, draft-katz-yeung-ospf-traffic-08.txt, Internet Draft, Work in Progress.

[RFC-2328] J.Moy, RFC 2328, ôOSPF Version 2ö, STD 54, IETF Standard Track, April 1998.

[RFC-2370] R.Coltun, RFC 2370, Standard Track, "The OSPF Opaque LSA Option", July 1998.

10. Acknowledgments

The authors would like to thank Bernard Sales, Emmanuel Desmet and Gert Grammel for their constructive comments, David Griffith for its active contribution and Bart Rousseau for its revision efforts.

11. Author’s Addresses

Dimitri Papadimitriou (Editor)
Francis Wellesplein, 1
B-2018 Antwerpen, Belgium
Phone: +32 3 240-8491
Email: dimitri.papadimitriou@alcatel.be

Fabrice Poppe (Alcatel)
Francis Wellesplein, 1
B-2018 Antwerpen, Belgium
Phone: +32 3 240-8006
Email: fabrice.poppe@alcatel.be

Jim Jones (Alcatel)
3400 W. Plano Parkway,
Plano, TX 75075, USA
Phone: +1 972 519-2744
Email: jim.d.jones@alcatel.com

Senthil Venkatachalam (OPNet)
Email: svenkatachalam@opnet.com

Sudheer Dharanikota (Consulting)
Email: sudheer@ieee.org

Raj Jain (Nayna Networks)
157 Topaz St.,
Milpitas, CA 95035, USA
Phone: +1 408 956-8000X309
Email: raj@nayna.com

Riad Hartani (Caspian Networks)
170 Baytech Drive,
San Jose, CA 95134, USA
Phone: +1 408 382-5216
Email: riad@caspiannetworks.com

Yong Xue (WorldCom)
Ashburn, VA, USA
Phone: +1 703 886-5358
Email: yong.xue@wcom.com

Appendix A: SRLG Computational Model

A.1 SRLG Concept and Example

The present model is intended to be used to automate the discovery of the Shared Risk Link Groups (SRLGs) at a given layer for a given physical resource type.

Note that a typical resource type can be a fiber, a fiber sub-segment, a fiber segment or a fiber trunk. SRLG definition and properties were described in Section 3.

Example:

The following example referring to Figure 5 (for the physical network topology) offers some clarification. Let assume that
- N1, N2, N3, and N4 represent locations that are linked by the fiber sub-segments,
- A, B, C, D and E be fiber segments,
- and F1 (ACD), F2 (AB), F3 (BCD) and F4 (DE) are fiber links routed over the fiber segment topology.

```
         N1            N2
         |             |
         |             |                            F1
         |A            |D                  N1 ------------ N2
         |             |                   |               |
         |             |                   |               |
         |      C      |                   |               |
```

```
         x-------------x                   |F2             |F4
         |             |                   |               |
         |             |                   |               |
         |             |                   |       F3      |
         |B            |E                  N3 ------------ N4
         |             |
         |             |
         N3            N4
```

Figure 4. Correlation between Fiber segment and Fiber link topology

In such a physical topology the obvious SRLGs are the following:
- {F1, F2} both going down when segment A breaks
- {F1, F3} both going down when segment C breaks
- {F1, F4} both going down when segment D breaks
- {F2, F3} both going down when segment B breaks
- {F3, F4} both going down when segment E breaks

These five SRLGs can be replaced by two SRLGs, S1 = {F1, F2, F3} and S2 = {F1, F3, F4}, where S1 and S2 constitute the minimum edge covering with cliques of the Shared Risk Relationship (SRR) graph that can be drawn between F1, F2, F3, F4 (see Figure 5). A clique of a graph G is a sub-graph of G in which every two nodes are connected by an edge. This decomposition is unique. If there was an additional dependency between F2 and F4, there would be a unique SRLG, S = {F1, F2, F3, F4}.

```
      F1 ------- F4
      |  \        |
      |   \       |
      |    \      |
      |     \     |
      |      \    |
      |       \   |
      |        \  |
      F2 ------- F3
```

Figure 5. SRR Graph between fiber link and segment

Although R1 = F1-F2-F3 and R2 = F4 are diverse path between locations N2 and N4 in the fiber topology (link and node disjointness), they are not diverse with respect to SRLGs. This is because both R1 and R2 cover SRLG S2, which contains F1, F3 (part of R1) and F4 (part of R2). SRLGs are thus a way of formalizing the propagation of (link-based) risk dependencies from server to client layers.

The rules guiding the definition of minimum set of SRLGs for more complex physical network topologies will be addressed in a future version of this document.

A.2 Risk Type

As specified up to now, the current approach considers that each of the network resource may experience one or more failure type(s). The same applies to geographical locations: a given location might experience one or more failure types. Moreover, by applying the SRLG properties, a network resource failure can potentially cover more than one geographical location. Consequently, some heuristics must be introduced to keep the SRLG computational complexity limited.

In order to limit the computational complexity, we define the following heuristics when considering the SRLG computation with respect to the type of risk:

1. The set of risk types associated to network resources corresponds exactly to the set of resource type failure.
   - So, for instance, the risk type associated to a fiber segment is a fiber segment failure. The same principle applies for other network resources such as fiber link, fiber sub-segment and fiber trunk. Consequently, the current approach does not consider a finest granularity for the network resource failure than the one referred by their type.

2. A risk type associated to a geographical structure covers exactly the region where it is defined. Moreover, a geographical failure is limited to a given location and does not contaminate the neighboring locations or generate another failure
   - For instance, we consider that an earthquake covers exactly one region or one area and that such a failure does not generate a hurricane impacting the neighboring locations. So, there is no correlation between geographical failures.

3. Each of the network resources covers exactly one geographical logical structure.
   - Consequently, when a geographical failure occurs, it generates a failure impacting the entire network resources included within the corresponding location. Hence, there is an ON/OFF relationship between geographical and network resource failures.

Consequently, when considering network resources, the risk type associated to an SRLG is defined as the potential failure of one (or more than one) instance belonging to a given resource type or the potential failure of any of its hierarchical resource dependence.

Thus, we define the concept of SRLG with respect to a given resource type and subsequently to the Risk Type (RT) to which this resource type refers. Moreover, for a given resource type (or class), each of the resources are identified by a Resource Identifier (RID). Since each of these resource classes corresponds to an SRLG class, we can assign an identifier to each of the SRLG classes members and define their value as the SRLG identifier.

Moreover, by applying the defined heuristics above, the SRLG identifiers can be grouped together by taking into account their geographical location. The latter is encoded by identifying the region identifier (region ID) including the resource identifiers to which the SRLG refers.

A.3 Calculation of Shared Risk Link Groups

In the calculation method, shared-risk(RID i, RID j, RT) is TRUE only if the resource identifiers RID i and RID j belong to the same SRLG with respect to the risk type RT. The risk types considered here are related the fiber segment, the fiber sub-segment and the fiber link risk failure.

A recursive calculation of shared-risk proceeds as follows:

```
shared-risk(RID i, RID j, RT) =
     at-risk(RID i, RT)
             and at-risk(RID j, RT)
             and (RID i = RID j
                     or (exists RID k, RID l
                     such that
                             depends-on(RID i, RID k)
                             and depends-on(RID j, RID l)
                             and shared-risk(RID k, RID l, RT)))
```

In this calculation:
- at-risk(RID i, RT) is TRUE only if RID is susceptible to a risk of type RT, either directly, or indirectly, through the failure of one of the resources it depends on.
- depends-on(RID i, RID j) is TRUE only if RID i fails as soon as RID j fails.

If we refer to the example detailed in section 1.1, then shared-risk(F1, F2, [fiber segment failure]) = TRUE because depends-on(F1, A) = TRUE, depends-on(F2, A) = TRUE and at-risk(A, [fiber segment failure]) = TRUE, the latter simply because A is a fiber segment.

A.4 Practical Method for SRLG Calculation

The recursive formula presented in the previous section does not directly lead to an efficient algorithm. ItÆs top-down nature illustrates nicely the recursive nature of the SRLG concept, but the calculation of the SRLGs in a top-down fashion would be totally inefficient, entailing the calculation of the same SRLGs in lower network layers over and over again.

A more efficient algorithm can be obtained from a bottom-up calculation. Figure 6 illustrates this by using the example we introduced in the Section A.1 and in by introducing the concept of Shared Risk Relationship Graph (SRR) which defines the membership of a resource belonging to the same SRLG.

```
          F1 ---------- F4
```

```
             |  \      ^   |
             |   \     |   |
             |    \    |   |        Fiber SRR Graph
         --->|     \   |   |<---
        |    |      \  |   |    |   where F1=ACD, F2=AB, F3=BCE, F4=DE
        |    |      ^\ |   |    |
        |    |      | \|   |    |
        |    |      |  |   |    |
        |    |      |  |\  |    |
        |    |      |  | \ |    |
        |    F2 ----|--|-- F3   |
        |        ^  |  |        |
        |        |  |  |        |
      +++++++++++++++++++++++++++++
        |        |  |  |        |
        |        |  |  |        |
         --- A   |  |   -- D    |
                 |  |           |
                 |  |           |
                 |  C           |   Fiber Segment SRR Graph
                 |              |
                 |              |
             B --          E ---
```

Figure 6. Bottom-up calculation of Shared Risk Relationships

For the calculation of a set of SRLGs, we need to calculate a Shared Risk Relationship (SRR) graph. The bottom-up calculation of the fiber SRR graph proceeds as follows:

- Step 1. For each fiber segment, there is an SRR between every two fibers contained in that segment (vertical arrows in Figure 6.)

- Step 2. For every SRR between two fiber segments, there is an SRR between every two fibers contained in either of the two fiber segments.

In the previous example, there are no SRRs between fiber segments, and the calculation stops after Step 1.

A.5 Application of the Model

The model is intended to be used to automate the discovery of the SRLGs at a given layer for a given risk type (RT).

The dependencies may be confined to one layer, e.g. the dependency of an optical link on a node (for instance, a DWDM system) to which it is connected, when the RT = [Node failure]. Dependencies may also extend over layer boundaries, e.g. the dependency of a TDM circuit (or sub-channel) in an SDH network established by using an optical channel (or wavelength) through the optical network that is the server layer of the SDH network, when RT = [fiber failure].

Let two optical network resources RID i and RID j within the same layer share a common risk of type RT. Let this risk type be tied to a lower layer, referred to as the risk layer. To enable this layer to infer shared-risk(RID i, RID j, RT), its server layer should advertise the following information:

shared-risk(component_1, component_2, RT)

where component_1 are services of the server layer on which RID i relies and component_2 are services of the serving layer on which RID j relies, respectively.

If the server layer is not the risk layer, the latter has to infer this knowledge from what its server layer is advertising. If shared risk relationships are not advertised, client layers should at least be able to query from their server layer the shared risk relationships between the services they receive.

Some dependencies do not lend themselves easily to automatic discovery. For instance, it is hardly imaginable that the process of finding out through which fiber segments a fiber goes can be automated. This means that part of the image of depends-on (RID i, RID j) will have to be provided æmanuallyÆ by the operator or be at least statically configured into a centralized repository.

More formally, an efficient calculation of shared risk link relationships relies on two things:

- In the lowest network layer with elements susceptible to the risk type RT that is considered, every network element RID j susceptible to the risk RT constitutes an SRR on its own, that is, (RID j, RID j) satisfies the recursive formula;

- Every SRR that has been discovered in one network layer leads to SRRs in the next higher network layer. In particular, two next higher layer network elements (RID i, RID j) depending on lower layer network elements that have an SRR satisfy the recursive formula. In order to allow an efficient calculation of the shared risk relationships in the next higher layer (e.g. the fiber layer), the shared risk relationships that were discovered in lower layers (e.g. the fiber segment layer) are stored in SRR graphs. This way, the recalculation of lower layer shared risk relationships can be avoided.

Full Copyright Statement