

CCAMP Working Group
Internet Draft
Intended status: Standard Track
Expires: August 17, 2013

Zafar Ali
George Swallow
Clarence Filsfils
Ori Gerstel
Stewart

Bryant

Matt

Hartley

Cisco

Systems

February 18,

2013

**Extended Encoding Scheme for Shared List Link Group (SRLG)
draft-ali-ccamp-extended-srlg-00.txt**

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on August 17, 2013.

Copyright Notice

Copyright (c) 2013 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF

Contributions published or made publicly available before November 10, [2008](#). **The person(s) controlling the copyright in some of this material** may not have granted the IETF Trust the right to allow modifications of

such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Abstract

SRLGs play a key role in routing resiliency and capacity planning of multi-domain and multi-layer networks. Notion of SRLG are used to select a backup path that is disjoint from the primary path, to ensure disjointness of circuits and to avoid catastrophic partitioning outages.

In the current specifications, SRLG is identified as a 32 bit number that is unique within an IGP domain [[RFC4202](#)]. There are many limitations to this approach of encoding SRLGs, especially in a multi-layer network. This draft outlines these limitations and suggests components of extended SRLG encoding scheme to address them.

Conventions used in this document

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

Table of Contents

Copyright Notice.....	1
1 . Introduction.....	3
2 . Requirements.....	3
2.1 . SRLG Scaling.....	3
2.2 . Global SRLG Identification.....	4
2.3 . Risk Management.....	4
3 . Components of Extended SRLG.....	5
3.1 . SRLG Filtration.....	5
3.2 . Globally Unique SRLGs.....	6
3.3 . SRLG Risk Management.....	6
4 . Extended SRLG Encoding.....	6
5 . Protocols extensions for extended SRLG Encoding.....	7
6 . Security Considerations.....	7
7 . IANA Considerations.....	7
8 . Acknowledgments.....	7
9 . References.....	7
9.1 . Normative References.....	7
9.2 . Informative References.....	8

1. Introduction

[RFC4202] defines notion of Shared Risk Link Group (SRLG). OSPF and IS-
IS extension for flooding SRLGs are defined in [[RFC4203](#)] and [[RFC5307](#)],
recording respectively. RSVP-TE signaling extensions for SRLG exclusion and
are defined in [[RFC4874](#)] and [[DRAFT-SRLG-RECORDING](#)], respectively.

In current specifications, SRLG is identified as a 32 bit number that
is unique within an IGP domain. There are many limitations to this
approach for encoding SRLGs, especially in multi-domain and multi-layer networks.

Section2 outlines these restrictions and states the associated requirements.

[Section 3](#) outlines components of the extended SRLG encoding format to address
these requirements. The extended SRLG encoding format and the associated
protocols extension(s) are intentionally left for a future version/ companion
document(s).

2. Requirements

The section outlines the requirements for extending SRLG beyond an IGP
domain scoped 32 bit number. Some of these requirements are also noted in
[DRAFT-SRLG-INFERENCE].

2.1. SRLG Scaling

A zealous operator could assign an SRLG for each risk, including (but
not limited to) a building, a floor of a building, a bridge, a side of a
rail-track, a side of a highway, and an amplifier. This operator could
easily have hundreds of SRLGs for Label Switch Paths (LSPs) transiting domain it
operates. Similarly, there are technologies (e.g., DWDM) where it is possible to
have hundreds of SRLGs associated with LSPs using such underlying
technology.

This presents a scaling issue with operations of the network, e.g.,

during constrained-based path computation of intra-domain LSPs. This also presents a scaling issue when such networks are used in multi-domain and multi-layer environment, e.g., in IP over DWDM network, there may be hundreds of SRLGs along a given IP/ MPLS link (inherited from underlying DWDM LSP). It may not scale for the IP/ MPLS layer to learn hundreds of SRLGs per link and flood them into its IGP database. This may impact flooding speed, topology database size and especially constrained-based path computation complexity and performance.

In the light of the above, finding mechanism(s) to scale SRLG is a requirement for GMPLS networks, especially in inter-domain/ inter-layer environment.

2.2. Global SRLG Identification

Currently SRLGs are defined and supported within domains. This limits the usefulness of SRLGs in an inter-domain environment, as elaborated in the following cases.

- There are cases where two different Service Providers (SPs) may be sharing the same fate (facility) for TE links within domains administrated by them. For example, if a client Service Provider SP-C leases two LSPs LSP1 and LSP2 respectively from two server SPs SP-S1 and SP-S2 who themselves lease fibers from the same submarine duct, the SP-C would like to know when the two LSPs LSP1 and LSP2 share the same submarine risk. With current definition of SRLGs this is not possible. This is because SP-S1 uses an SRLG numbering completely independent from SP-S2. For example, SP-S1 might identify the submarine risk as SRLG23 while SP-S2 identifies it as SRLG47. Even if client SP (SP-C) is able to discover SRLGs along LSP1 and LSP2 (e.g., using SRLG recording [[DRAFT-SRLG-RECORDING](#)] SP-C learns that the LSP1 is exposed to SRLG23 and LSP2 exposed to SRLG47), the SP-C problem is not resolved: it has no way to know that LSP1 and LSP2 are actually sharing the same risk.
- If SP-C in the above example would like to request LSP2 to be SRLG diverse from LSP1 using SRLG recording [[DRAFT-SRLG-RECORDING](#)] and SRLG XRO/EXRS [[RFC4874](#)], there is no guarantee that LSP2 route is SRLG diverse. This is because knowing that LSP1 is exposed to SRLG23, SP-C cannot realize a path from SP-S2 which is disjoint from SRLG23 of SP-S1 (as SRLG23 means something else for SP-S2). Similarly, SRLG inclusion also does not work using the current SRLG encoding scheme.

- At present, SRLG administration is completely up to the SPs.
Therefore, SRLG values in an inter-domain environment may collide. Considering the above example, SP-S1 may have assigned SRLG12 to a resource used by LSP1, whereas client SP-C may already be using SRLG12 to identify a different resource in its network. Even though these two resources may not share any risk, they are not SRLG diverse (assumed to share risk in GMPLS control plane).

In the light of the above, finding mechanism(s) to maintain consistency of SRLG in an inter-domain environment is a requirement.

2.3. Risk Management

Not all resources in a network have same level of availability. Some resources are more prone to failures, e.g., a fiber trunk running close to utility wires is more likely to suffer from accidental cuts than a fiber trunk running in isolation. Metro links are more susceptible to cuts than rural links, and aerial fiber is susceptible to storm induced outages.

Consider the example where a client Service Provider (SP) SP-C leases LSP1 from server SP (SP-S1). Availability of LSP1 is typically part of service level agreement (SLA) between SP-C and SP-S1, e.g., SP-C may request 99.999% (five-nines) of availability. Given high availability requirement for LSP1, SP-S1 needs to route LSP1 such that it uses resources with better than 99.999% availability. Furthermore, given a set of underlying resources, SP1 should also be able to estimate availability of LSP1 connection. How availability of a connection given availability of its underlying resources is estimated is beyond the scope of this document but, if availability is represented as a number between 0 and 1, a multiply function can be used for this calculation.

In the light of the above, finding mechanism(s) to quantify risk associated with a resource is a requirement.

3. Components of Extended SRLG

This section outline components that form basis for extended SRLG encoding scheme.

3.1. SRLG Filtration

A way to address SRLG scaling requirement mentioned in [Section 2](#) is to associate a priority field to the SRLG and use it as a mechanism to observe SRLG filtering. For example, in a multi-layer network, only higher priority SRLGs may be requested or exposed to the client layer. Alternatively, the client may request all the SRLG's from the server and store them locally in its SRLG database but only flood in its client control-plane (ISIS, OSPF) the more important (higher priority) SRLG's.

Examining a resource type associated with an SRLG may also be used to filter

SRLG information in multi-domain/ multi-layer networks. E.g., SP-S1 may not export SRLGs of amplifiers used along path of LSP1 to the client SP (SP-C) running IP services. This document suggests characterization of the following resource types:

- Optical section: A fiber that connects two optical NEs(e.g., amplifiers). Also termed OTS in ITU parlance.
- Optical line: A fiber that connects two optical switching elements (e.g., ROADMs). Also termed OMS in ITU parlance.
- Optical path: an optical connection that connects two client ports, e.g., a port P1 on node N1 to a port P2 on node N2. Also termed Och Connection in ITU parlance.
- Fiber Duct: Conduit carrying fibers (which represent optical sections).

- Building: Building hosting multiple network elements, and represents a common risk, e.g., during a terror attack. Such a building may host both transport and client gear.
- Optical NE: Amplifier, ROADM or other optical NE used along an optical TE link.
- Power feed: a common power source feeding multiple NEs
- Geographic region: an area susceptible to a disaster such as earthquake or flood.
- More may be added in future revision(s).

3.2. Globally Unique SRLGs

A way to address global uniqueness of the SRLG is to associate Autonomous System (AS) number of the AS that originated the SRLG value.

3.3. SRLG Risk Management

Risk management requirements are discussed in [section 2](#). As SRLGs are used to characterize resources, associating a resource availability field to the SRLG can satisfy risk management requirements. Specifically, availability of a resource as defined in the following can be used for this purpose.

Availability = $MTBF / (MTBF + MTTR)$, where

MTBF is mean time between failures of the resource,

MTTR is mean time to repair a failure of the resource.

4. Extended SRLG Encoding

Motivation for this version is to discuss components of SRLG and hence the extended SRLG encoding is intentionally left for a future revision. Nonetheless the idea is to augment the currently defined 32-bit Shared Risk Link Group Value with the following parameters:

- o SRLG priority.
- o Type of the SRLG resource.
- o Availability of the SRLG SRLG resource.
- o SRLG Originator's AS Number.

5. Protocols extensions for extended SRLG Encoding

Extension(s) for protocols that use SRLG values for their operations (e.g., OSPF, ISIS, RSVP-TE, PCE, etc.) are to be added in future version of this document or companion document(s).

6. Security Considerations

Security considerations related to the extended-SRLG encoding are left for future revision of the document.

7. IANA Considerations

This version of the document does not require any IANA considerations.

8. Acknowledgments

Authors would like to acknowledge valuable feedback from many service providers. Individual acknowledgements will be added in future version of the document.

9. References

9.1. Normative References

- [RFC4202] Kompella, K., Ed., and Y. Rekhter, Ed., "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4202](#), October 2005.
- [RFC4203] Kompella, K., Ed., and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4203](#), October 2005.
- [RFC5307] Kompella, K., Ed., and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 5307](#), October 2008.
- [RFC4874] Lee, CY., Farrel, A., and S. De Cnodder, "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", [RFC 4874](#), April 2007.
- [DRAFT-SRLG-RECORDING] F. Zhang, D. Li, O. Gonzalez de Dios, C. Margaria, M. Hartley, Z. Ali, "RSVP-TE Extensions for Collecting SRLG Information", [draft-ietf-ccamp-rsvp-te-srlg-collect.txt](#), work in progress.

9.2. Informative References

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

Authors' Addresses

Zafar Ali
Cisco Systems
Email: zali@cisco.com

George Swallow
Cisco Systems
swallow@cisco.com

Clarence Filsfils
Cisco Systems
cfilsfil@cisco.com

Ori Gerstel
Cisco Systems
Email: ogerstel@cisco.com

Stewart Bryant
Cisco Systems
stbryant@cisco.com

Matt Hartley
Cisco Systems
Email: mhartley@cisco.com

