

Network Working Group
Internet-Draft
Intended status: Informational
Expires: November 13, 2014

Pierre Francois
IMDEA Networks Institute
Clarence Filsfils
Cisco Systems, Inc.
Bruno Decraene
Orange
Rob Shakir
BT
May 12, 2014

Use-cases for Resiliency in SPRING
draft-ietf-spring-resiliency-use-cases-00

Abstract

This document describes the use cases for resiliency in SPRING networks.

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 November 13, 2014.

Copyright Notice

Copyright (c) 2014 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.

Table of Contents

- [1. Introduction](#) [3](#)
- [2. Path protection](#) [4](#)
- [3. Management free local protection](#) [4](#)
 - [3.1. Management free bypass protection](#) [5](#)
 - [3.2. Management-free shortest path based protection](#) [5](#)
- [4. Managed local protection](#) [6](#)
 - [4.1. Managed bypass protection](#) [6](#)
 - [4.2. Managed shortest path protection](#) [6](#)
- [5. Co-existence](#) [7](#)
- [6. References](#) [7](#)
- [Authors' Addresses](#) [7](#)

1. Introduction

SPRING aims at providing a network architecture supporting services with tight SLA guarantees [1]. This document reviews various use cases for the protection of services in a SPRING network. Note that these use cases are in particular applicable to existing LDP based and pure IP networks.

Three key alternatives are described: path protection, local protection without operator management and local protection with operator management.

Path protection lets the ingress node be in charge of the failure recovery, as discussed in Section 2.

The rest of the document focuses on approaches where protection is performed by the node adjacent to the failed component, commonly referred to as local protection techniques or Fast Reroute techniques.

We discuss two different approaches to provide unmanaged local protection, namely link/node bypass protection and shortest path based protection, in Section 3.

A case is then made to allow the operator to manage the local protection behavior in order to accommodate specific policies, in Section 4.

The purpose of this document is to illustrate the different approaches and explain how an operator could combine them in the same network (see Section 5). Solutions are not defined in this document.




Figure 1: Reference topology

We use Figure 1 as a reference topology throughout the document. All link metrics are equal to 1, with the exception of the links from/to A and Z, which are configured with a metric of 100.

2. Path protection

A first protection strategy consists in excluding any local repair but instead use end-to-end path protection.

For example, a Pseudo Wire (PW) from A to Z can be "path protected" in the direction A to Z in the following manner: the operator configures two SPRING paths T1 and T2 from A to Z. The two paths are installed in the forwarding plane of A and hence are ready to forward packets. The two paths are made disjoint using the SPRING architecture.

T1 is established over path {AB, BC, CD, DE, EZ} and T2 over path {AF, FG, GH, HI, IZ}. When T1 is up, the packets of the PW are sent on T1. When T1 fails, the packets of the PW are sent on T2. When T1 comes back up, the operator either allows for an automated reversion of the traffic onto T1 or selects an operator-driven reversion. The solution to detect the end-to-end liveness of the path is out of the scope of this document.

From a SPRING viewpoint, we would like to highlight the following requirement: the two configured paths T1 and T2 MUST NOT benefit from local protection.

3. Management free local protection

This section describes two alternatives to provide local protection without requiring operator management, namely bypass protection and shortest-path based protection.

For example, a demand from A to Z, transported over the shortest paths provided by the SPRING architecture, benefits from management-free local protection by having each node along the path automatically pre-compute and pre-install a backup path for the destination Z. Upon local detection of the failure, the traffic is repaired over the backup path in sub-50msec.

The backup path computation should support the following requirements:

- o 100% link, node, and SRLG protection in any topology
- o Automated computation by the IGP
- o Selection of the backup path such as to minimize the chance for transient congestion and/or delay during the protection period, as reflected by the IGP metric configuration in the network.

3.1. Management free bypass protection

One way to provide local repair is to enforce a failover along the shortest path around the failed component, ending at the protected nexthop, so as to bypass the failed component and re-join the pre-convergence path at the nexthop. In the case of node protection, such bypass ends at the next-nexthop.

In our example, C protects Z, that it initially reaches via CD, by enforcing the traffic over the bypass {CH, HD}. The resulting end-to-end path between A and Z, upon recovery against the failure of C-D, is depicted in Figure 2.




Figure 2: Bypass protection around link C-D

3.2. Management-free shortest path based protection

An alternative protection strategy consists in management-free local protection, aiming at providing a repair for the destination based on shortest path state for that destination.

In our example, C protects Z, that it initially reaches via CD, by enforcing the traffic over its shortest path to Z, considering the failure of the protected component. The resulting end-to-end path between A and Z, upon recovery against the failure of C-D, is depicted in Figure 3.




Figure 3: Reference topology

4. Managed local protection

There may be cases where a management free repair does not fit the policy of the operator. For example, in our illustration, the operator may want to not have C-D and C-H used to protect each other, in fear of a shared risk among the two links.

In this context, the protection mechanism must support the explicit configuration of the backup path either under the form of high-level constraints (end at the next-hop, end at the next-next-hop, minimize this metric, avoid this SRLG...) or under the form of an explicit path.

We discuss such aspects for both bypass and shortest path based protection schemes.

4.1. Managed bypass protection

Let us illustrate the case using our reference example. For the demand from A to B, the operator does not want to use the shortest failover path to the nexthop, {CH, HD}, but rather the path {CG,GH,HD}, as illustrated in Figure 4.




Figure 4: Managed bypass protection

4.2. Managed shortest path protection

In the case of shortest path protection, the case is the one of an operator who does not want to use the shortest failover via link C-H, but rather reach H via {CG, GH}.

The resulting end-to-end path upon activation of the protection is illustrated in Figure 5.




Figure 5: Managed shortest path protection

5. Co-existence

The operator may want to support several very-different services on the same packet-switching infrastructure. As a result, the SPRING architecture SHOULD allow for the co-existence of the different use cases listed in this document, in the same network.

Let us illustrate this with the following example.

- o Flow F1 is supported over path {C, C-D, E}
- o Flow F2 is supported over path {C, C-D, I}
- o Flow F3 is supported over path {C, C-D, Z}
- o Flow F4 is supported over path {C, C-D, Z}
- o It should be possible for the operator to configure the network to achieve path protection for F1, management free shortest path local protection for F2, managed protection over path {C-G, G-H, Z} for F3, and management free bypass protection for F4.

6. References

[1] Filsfils, C., Previdi, S., Bashandy, A., Decraene, B., Litkowski, S., Horneffer, M., Milojevic, I., Shakir, R., Ytti, S., Henderickx, W., Tantsura, J., and E. Crabbe, "Segment Routing Architecture", [draft-filsfils-rtgwg-segment-routing-01](#) (work in progress), October 2013.

Authors' Addresses

Pierre Francois
 IMDEA Networks Institute
 Leganes
 ES

Email: pierre.francois@imdea.org

Clarence Filsfils
Cisco Systems, Inc.
Brussels
BE

Email: cfilsfil@cisco.com

Bruno Decraene
Orange
Issy-les-Moulineaux
FR

Email: bruno.decraene@orange.com

Rob Shakir
BT
London
UK

Email: rob.shakir@bt.com