

TEAS Working Group
Internet-Draft
Intended status: Informational
Expires: 28 August 2026

I. Busi
Huawei Technologies
J.-F. Bouquier
Vodafone
F. Peruzzini
FiberCop
P. Volpato
Huawei Technologies
P. Manna
Cisco
24 February 2026

Applicability of Abstraction and Control of Traffic Engineered Networks
(ACTN) for Packet Optical Integration (POI) service assurance
draft-ietf-teas-actn-poi-assurance-04

Abstract

This document extends the analysis of the applicability of Abstraction and Control of TE Networks (ACTN) architecture to Packet Optical Integration (POI) to cover multi-layer service assurance scenarios. Specifically, the ACTN architecture enables the detection and handling of different failures that may happen either at the optical or the packet layer. It is assumed that the underlying transport optical network carries end-to-end IP services such as L2VPN or L3VPN connectivity services, with specific Service Level Agreement (SLA) requirements.

Existing IETF protocols and data models are identified for each multi-layer (packet over optical) service assurance scenario with a specific focus on the MPI (Multi-Domain Service Coordinator to Provisioning Network Controllers Interface) in the ACTN architecture.

About This Document

This note is to be removed before publishing as an RFC.

The latest revision of this draft can be found at <https://IETF-TEAS-WG.github.io/actn-poi-assurance/draft-ietf-teas-actn-poi-assurance.html>. Status information for this document may be found at <https://datatracker.ietf.org/doc/draft-ietf-teas-actn-poi-assurance/>.

Discussion of this document takes place on the Traffic Engineering Architecture and Signaling Working Group mailing list (<mailto:teas@ietf.org>), which is archived at <https://mailarchive.ietf.org/arch/browse/teas/>. Subscribe at <https://www.ietf.org/mailman/listinfo/teas/>.

Source for this draft and an issue tracker can be found at
<https://github.com/IETF-TEAS-WG/actn-poi-assurance>.

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 <https://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 28 August 2026.

Copyright Notice

Copyright (c) 2026 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
2. Conventions and Definitions	5
2.1. Terminology	5
3. Reference Network Architecture	6
3.1. Reference Network	9
4. YANG Data Models for the MPIs	10
5. Multi-layer Fault Management	11
5.1. Reference scenario for multi-layer faults	11
5.2. Optical Network Failures	12
5.3. Cross-layer Link Failures	14
5.4. Router Node Failures	16
5.5. VPN Service Fault Monitoring	16

5.5.1. Remedial Action and Resolution	17
6. Multi-layer Performance Management	17
6.1. Optical performance management	18
6.2. End-to-end IP performance management	19
7. Multi-layer Resiliency	19
7.1. Optical Network Failures	20
7.1.1. Optical restoration	21
7.1.2. Optical protection	23
7.2. Optical Network Maintenance	24
7.3. Cross-layer Link Failures	27
7.4. Router Node Failures	30
7.5. Multi-layer hitless reversion	32
8. Conclusions	34
9. Security Considerations	34
10. IANA Considerations	35
11. References	35
11.1. Normative References	35
11.2. Informative References	37
Acknowledgments	37
Contributors	37
Authors' Addresses	38

1. Introduction

Service assurance is a critical aspect of Operations, Administration and Management (OAM). It consists of activities and processes intended to guarantee a specified Service Level Agreement (SLA) for the customer of a telecommunication service. Service assurance includes both fault management, to correct service anomalies and network faults, and performance management, to monitor service and network parameters and provide early warning of potential service-related issues.

Within the scope of this document, service assurance is discussed in the context of a multi-layer, multi-domain network. This document leverages the Abstraction and Control of TE Networks (ACTN) framework [RFC8453] and further expands the analysis of its applicability to multi-layer packet-optical integrated networks [I-D.ietf-teas-actn-poi-applicability], adding considerations specific to fault and performance management scenarios.

As already highlighted in [I-D.ietf-teas-actn-poi-applicability], a multi-layer network is composed of an IP layer and an optical transport layer. A multi-domain network is composed of at least two different administrative domains (e.g. core and edge) under the control of the same organization (e.g. the same network operator). Service assurance applies to end-to-end L2VPN or L3VPN connectivity services configured over underlying transport optical paths that require multi-layer coordination.

To guarantee the SLAs associated with VPN services, service assurance is performed through collaboration among the ACTN control entities [RFC8453]: the Multi-Domain Service Coordinator (MDSC), acting as the top-level controller, and the Provisioning Network Controllers (PNCs) deployed in both the packet (P-PNC) and optical (O-PNC) layers.

This document aligns with current field operational procedures in optical networks and assumes that the O-PNC provides the MDSC with enough information for Root Cause Analysis (RCA), correlating an event/alarm related to an optical-network failure with impacted services at the IP layer.

The set of information shared by the O-PNC to the MDSC depends on local configuration adopted at the MDSC-PNC Interface (MPI) [RFC8453]. In general, this may include information about the optical path, tunnel, or fiber where the failure happened, together with its location and operational state (e.g., "down"), while hiding further topology details. This data is sufficient for the MDSC to perform multi-layer correlation and discover which IP links, Label Switched Paths (LSPs), and VPNs are affected.

The analysis of the YANG data models applicable to service assurance (fault and performance) is in scope of this document. The development of new YANG models/modules to support the missing functions is instead not in scope of the present document. To this extent, this document acts as a framework that provides a gap analysis and identifies topics for future work to be addressed in other documents.

A related effort is described in [I-D.ietf-ccamp-actn-optical-transport-mgmt], which enhances the ACTN architecture for optical networks with FCAPS (Fault, Configuration, Accounting, Performance, and Security) management capabilities, including the addition of rich-detail network management (RDNM) to the MPI. While that document addresses the general integration of YANG-based FCAPS capabilities into ACTN for optical networks, the present document focuses on multi-layer (packet and optical) service assurance in Packet Optical Integration (POI) scenarios, specifically analyzing fault detection, performance monitoring, and resiliency coordination across the packet and optical layers.

The document has the following organization. Section 2 lists the conventions and definitions used in the text. Section 3 discusses the reference network in scope for the relevant service assurance cases. Section 4 identifies the YANG data models applicable to service assurance and provides a gap analysis for modules that are still missing.

Section 5 identifies the possible faults, either in the optical layer, the IP layer, or both. Section 6 deals with performance management aspects of service assurance in a packet-optical integrated network. Finally, Section 7 discusses the protection mechanisms available for the most typical fault scenarios of a multi-layer, multi-domain network.

For each multi-technology scenario, the document analyzes how to use the interfaces and the data models of the ACTN architecture.

A summary of the gaps identified in this analysis is provided in Section 8.

Understanding the degree of standardization and the identified gaps will help assess the feasibility of integration between packet and optical Dense Wavelength Division Multiplexing (DWDM) domains (and optionally the Optical Transport Network (OTN) layer) from an end-to-end, multi-vendor service assurance perspective.

2. Conventions and Definitions

2.1. Terminology

TCA (Threshold Crossing Alert) A notification generated by a network

element or controller when a monitored performance parameter crosses a pre-configured threshold value. TCAs are used to signal potential SLA violations or network degradation conditions to higher-layer management systems. In the context of this document, TCAs are generated at the router or PNC level and propagated upward through the MDSC to the Orchestrator layer.

3. Reference Network Architecture

This document analyses several scenarios for service assurance in Packet and Optical Integration (POI) in which ACTN hierarchy is deployed to control a multi-layer and multi-domain network with two optical domains and two packet domains, as shown in Figure 1 of [I-D.ietf-teas-actn-poi-applicability], which is copied in Figure 1 below.

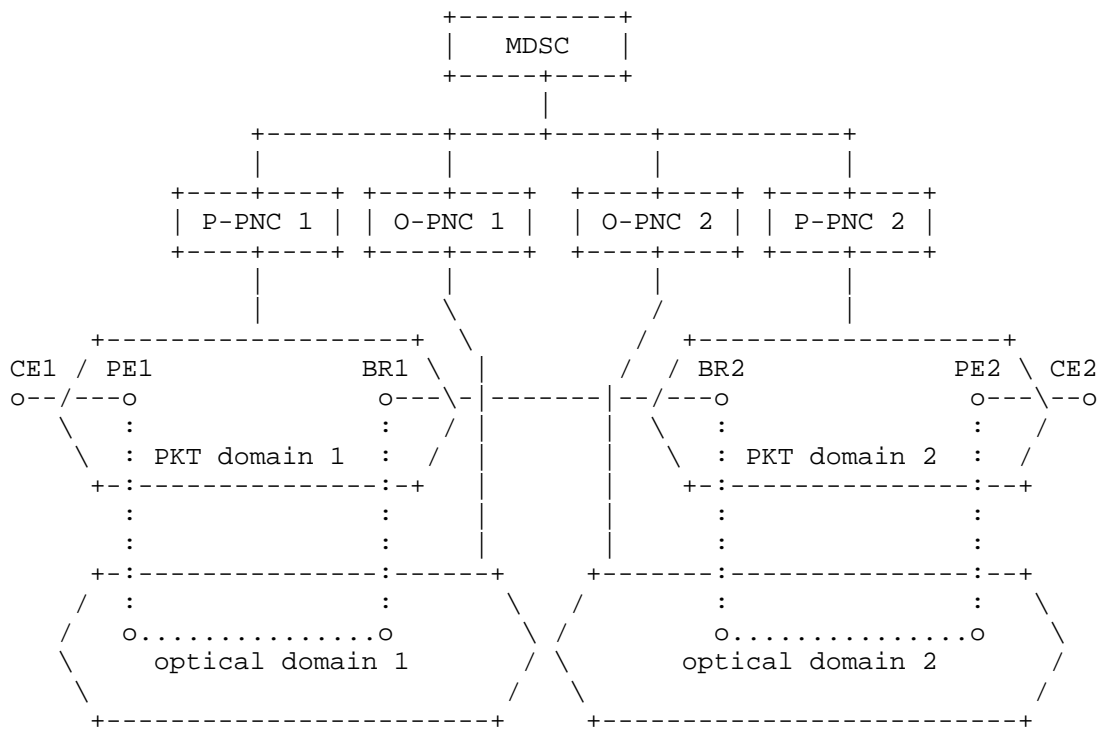


Figure 1: Reference Network (copy of Figure 1 of RFC YYYY)

EDITORS NOTE: Replace RFC YYYY with the RFC number of [I-D.ietf-teas-actn-poi-applicability] once it has been published.

In general, service assurance involves fault detection and localization; performance monitoring as well as re-routing (protection).

Two cases will be considered:

1. using grey interfaces on routers' ports, as outlined in [I-D.ietf-teas-actn-poi-applicability]
2. using colored optical interfaces on routers' ports, as outlined in [I-D.mix-teas-actn-poi-extension]

NOTE: It is not fully clear how many commonalities there are in service assurance for these two cases. This draft will start addressing both cases. At a later stage it will be assessed whether it is worthwhile keeping everything in a single draft or to split into two drafts.

The MDSC is responsible for coordinating the whole multi-domain, multi-layer (packet and optical) network. MDSC interacts with different Provisioning Network Controllers (O/P-PNCs) through the MPI interface. The MPI interface presents an abstracted topology to MDSC, hiding the technology-specific aspects of the network and the topology details (depending on the policy chosen regarding the level of abstraction supported).

Following the assumptions of section 2.1.2 of [I-D.ietf-teas-actn-poi-applicability], this document analyses scenarios where the MDSC uses the partial summarization approach to coordinate multi-domain/multi-layer path computation.

This document assumes that the basic configuration of the L2/L3 VPN services, and in particular of the multi-domain TE path in the packet layer, is the same as described in section 5 of [I-D.ietf-teas-actn-poi-applicability]. The only exception is that section Section 7 defines additional multi-layer mechanisms besides the local protection in the packet layer described in section 5.2.3 of [I-D.ietf-teas-actn-poi-applicability].

In this approach, the MDSC has complete visibility of the TE topology of the packet network domains and an abstracted view of the TE topology of the optical network domains. That means the MDSC has the capability of performing multi-domain/single-layer path computation for the packet layer. The MDSC needs to delegate the O-PNCs to perform local path computation within their respective domains. It uses the information received by the O-PNCs and its TE topology view of the multi-domain packet layer to perform multi-layer/multi-domain path computation.

P-PNCs are responsible for setting up the TE paths between any two PEs or BRs in their respective controlled domains, as requested by MDSC, and providing topology information to the MDSC.

O-PNCs are responsible for providing to the MDSC an abstract TE topology view of their underlying optical network resources. They perform single-domain local path computation, when requested by the MDSC. They also perform optical tunnel setup, when requested by the MDSC.

No GMPLS-UNI interaction between IP and Optical equipment is considered. This is also the assumption followed in this document: the MDSC performs the function of multi-layer/multi-domain path computation through the same mechanisms described in [I-D.ietf-teas-actn-poi-applicability].

TO DO: Complete the description of the pre-requisites of MDSC in the cases discussed.

The following list summarizes the main assumptions about how MDSC can handle the service assurance cases described in this document. Most of them have already been described in [I-D.ietf-teas-actn-poi-applicability].

1. MDSC has acquired all the topology and status information of both the IP and optical layers.
2. MDSC is fully aware of any multi-layer connections between the IP and the optical layers. It is also aware of the multi-domain interconnection links between different IP domains.
3. MDSC is aware of any topology or resource utilization change obtained in real time through coordination with the O/P-PNCs. This applies in the case of a fault or a maintenance activity involving either the IP or the DWDM layer.
4. MDSC coordinates the IP and DWDM protections and, as a result, the re-routing of traffic at both the IP and DWDM layer.
5. Before planned maintenance operation at the DWDM layer, MDSC instructs the P-PNC to move the affected IP traffic to another link in a hitless way. This is done before the event takes place. MDSC also coordinates with P-PNC to revert back the traffic on the original path when the maintenance event is concluded.

6. When the O-PNC detects a degradation of optical performance (e.g. a Threshold Crossing Alert (TCA) on PRE-FEC BER values sustained over a certain period of time), it alerts the MDSC so that the MDSC relates the warning to an IP link.
7. MDSC distinguishes between IP and Optical failures. For example, in the case of the failure of an IP port of a router, the IP traffic may be switched to a stand-by port, reusing the same Reconfigurable Optical Add-Drop Multiplexer (ROADM) optical resources (lambda, optical path) and keeping the end-to-end IP connection. If a remote IP node fails, then a re-route of optical resources takes place together with a switch of the local IP port in order to establish a new connection with a different IP node used for protection.

3.1. Reference Network

The following network topology will be considered to analyze and discuss the scenarios in Section 7.

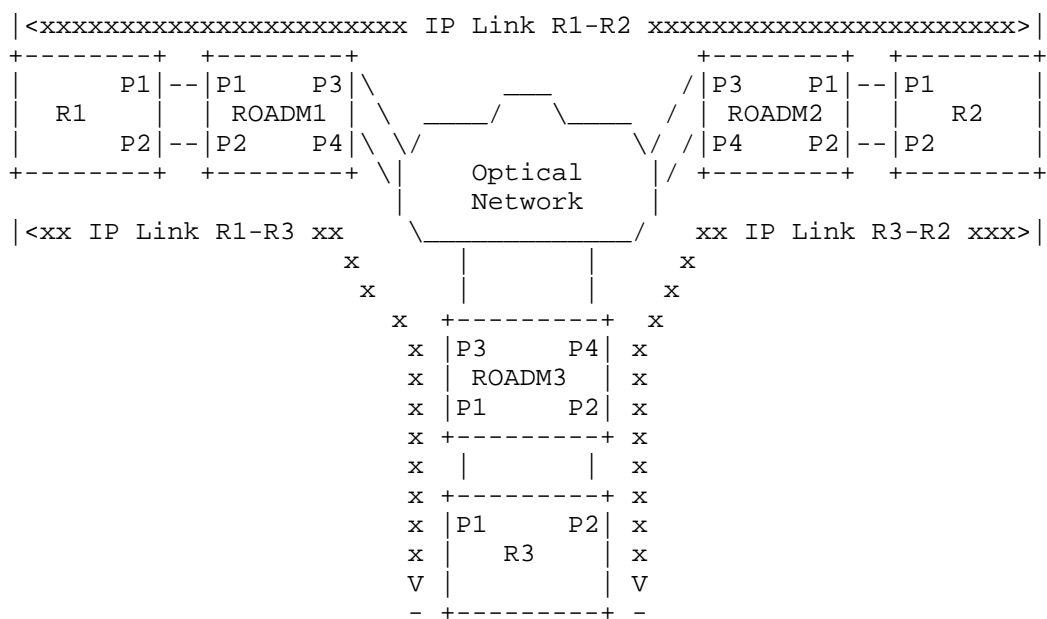


Figure 2: Reference Network

The network consists of three Points of Presence (POPs) geographically distributed. It is assumed that every POP hosts a Router (R1, R2, and R3 respectively) connected to a ROADM (ROADM1, ROADM2, and ROADM3). All the routers connect to their co-located

ROADMs with two Ethernet links (e.g. 100GE) for redundancy. In their normal operations, the routers may employ any local policy for traffic steering. For the scope of this document, it is assumed that the path that R1 uses to steer the IP traffic to R2 goes from port P1 of R1 to port P1 of R2 (thus going through port P1 of R1, ports P1 and P3 of ROADM1, ports P3 and P1 of ROADM2, port P1 of R2). R1 uses port P2 to steer the traffic to R3 instead. The IP link between R1 and R3 carries the IP services that are directed to R3 and is used by R1 as a detour path (backup path) to reach R2 if a failure occurs in the primary path across ROADM1 and ROADM2. The detour path also includes a second leg from R3 to R2. The detour path from R1 to R2, then includes: port P2 of R1, ports P2 and P4 of ROADM1, ports P3 and P1 of ROADM3, ports P1 and P2 of R3, ports P2 and P4 of ROADM3, ports P4 and P2 of ROADM2, and port P2 of R2. The connection between all ROADMs is based on two fibers. The optical paths all cross an optical network. For the scope of this document, it is assumed that some coordination mechanisms are employed at the optical layer so that when a failure happens on an optical path (for example, between ROADM1 and ROADM2), an optical backup path is activated. The mechanisms are assumed to be coordinated by O-PNC and MDSC, even if other methods may be also considered (e.g. Generalized MPLS (G-MPLS) based). Further details are given in the use cases described in Section 7.

4. YANG Data Models for the MPIs

The analysis of the data models potentially of interest for this document is still on-going. The set of YANG models identified so far includes the following items:

- * ietf-alarms defined in [RFC8632]
- * ietf-performance-monitoring defined in [I-D.yu-performance-monitoring-yang]
- * A YANG Data Model for Service Assurance [RFC9418]
- * A YANG Data Model for Network and VPN Service Performance Monitoring [RFC9375]
- * YANG Data Model for Layer-2 Network Model (L2NM) [RFC9291]
- * A YANG Data Model for Layer-3 Network Model (L3NM) [RFC9182]

The list will be progressively updated as the document evolves.

5. Multi-layer Fault Management

This section deals with the actions taken by the MDSC and the PNCs at the IP and optical layers to handle the occurrence of a failure in a multi-layer network. This set of actions is referred to as fault management and consists of steps such as fault detection, fault localization, and fault recovery. Specifically, this section analyzes failure detection and localization, while Section 7 further details fault recovery mechanisms. Depending on the point where a failure occurs, three use cases are considered: 1. The failure occurs in the optical layer, for example a fiber cut that triggers a Loss of Signal (LOS) alarm. This is discussed in section Section 5.2. 2. The failure occurs at the connection between a router and a ROADM (cross-layer link). Such a case is analyzed in section Section 5.3. 3. The failure occurs in the IP layer, for example a router experiences a hardware failure on a port that connects to its optical counterpart. This case is discussed in section Section 5.4.

5.1. Reference scenario for multi-layer faults

The following figure illustrates the reference scenario used to discuss the fault management cases.

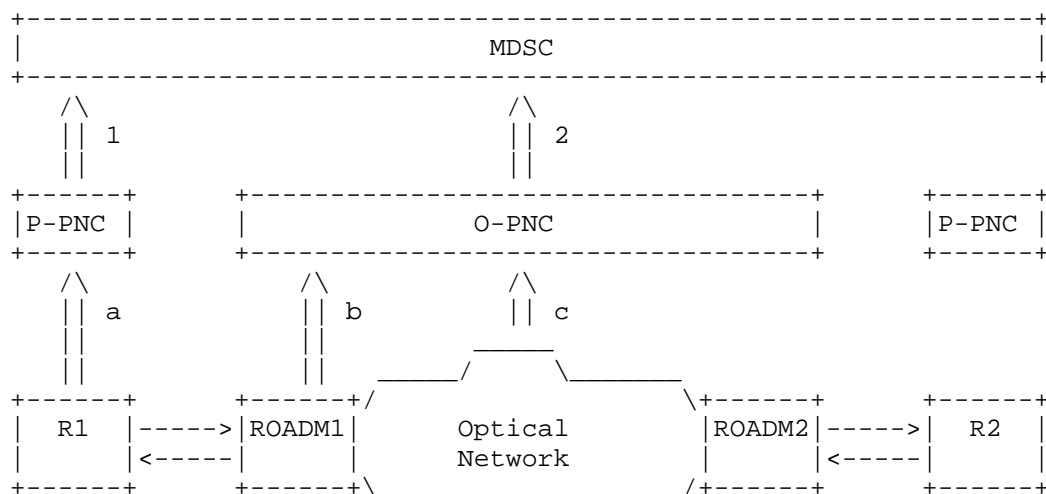


Figure 3: Reference scenario for multi-layer fault management

The MDSC is responsible for the correlation of the events. The notification about a failure (alarm, state change, etc.) is sent from the P-PNC (upstream arrow labelled "1" in the figure) and/or the O-PNC (upstream arrow labelled "2"), depending on the case

considered. It should be noted that only a single P-PNC is present in the network. The second box marked with the label "P-PNC" is represented only to simplify the schematics. Actually the P-PNC on the left and the P-PNC on the right are the same element.

In case of a failure in the IP layer, the router that detects it sends a corresponding notification message to the P-PNC. This is represented by the upstream arrow labelled "a". Similarly, a failure in the optical layer can be notified through messages sent by a ROADM (upstream arrow "b") or by a node within the optical core network (upstream arrow "c"). Again, depending on the specific case, multiple messages can be sent by IP and/or optical nodes to the corresponding PNC.

For simplicity, a router is connected to a ROADM via two unidirectional fibers, represented by the two arrows between them. ROADM 1 and ROADM 2 are considered to be the edge nodes of a larger optical core that may include several other components. The two connections between a router and a ROADM carry, in addition to data traffic, the signaling messages generated by the physical transmission layer, for example Local Failure Indication (LFI) or Remote Failure Indication (RFI). These messages provide supplementary information that an IP or an optical node may consider for failure detection and for providing further details in the upstream notification to a PNC.

5.2. Optical Network Failures

In this case, the O-PNC is fully responsible for the fault management (including failure detection, location and repair) within the optical domain.

The detailed mechanisms used by the O-PNC for intra-domain fault management are outside the scope of this document. Optical data plane standards provide a comprehensive set of OAM tools, defined in [ITU-T_G.709] and [ITU-T_G.798], that would assist O-PNC fault management, as described in [ITU-T_G.7710] and [ITU-T_G.874].

It is worth noting that the OAM tools, defined in [ITU-T_G.709] and [ITU-T_G.798], are fully standardized for the ODU, OTU and FlexO sub-layers but only functionally standardized for the optical media layer (i.e., OCh and OTSiA). This is not an issue since it is assumed that the optical NEs and O-PNC within a single domain are single-vendor.

However, the level of standardization of the OAM tools management requirements is sufficient to define standard requirements and data models at the MPI for multi-vendor, multi-domain and multi-layer fault management.

Even if in this case the fault management is fully under the responsibility of the O-PNC, it is still needed to inform the MDSC that there is a failure within the optical domain and that the O-PNC is working on it.

A failure within the optical network can cause secondary failures on multiple optical tunnels, which can in turn cause failures on the multi-layer IP links and on the L2VPN and L3VPN services whose traffic is sent over the failed tunnels.

For example, with a reference to Figure 7 of [I-D.ietf-teas-actn-poi-applicability], a failure within the optical network can cause a failure on the optical tunnel between NE11 and NE12. As a consequence, also the IP link between PE13 and BR11 is failed and the L2VPN/L3VPN xxx is also affected.

The O-PNC can report the operational status of the optical tunnels to the MDSC to let the MDSC know that the optical tunnel is down. The MDSC can then correlate the failure of the optical tunnel (e.g., the optical tunnel between NE11 and NE12 in Figure 7 of [I-D.ietf-teas-actn-poi-applicability]) with the secondary failures on the L2VPN/L3VPN whose traffic has been routed through that optical tunnel.

Comment: Need to discuss here why reporting the operational status of the optical tunnel is not sufficient to motivate the need for a more enhanced incident management as proposed in [I-D.feng-opsawg-incident-management]

The MDSC should also inform the OSS/orchestration layer about the failures on the affected L2VPN/L3VPN services through mechanisms that are outside the scope of this document.

Comment: Need further discussion about the behavior of P-PNC. The P-PNC can also discover that the multi-layer IP link is down (e.g., using BFD). However, I think that the fault management process in P-PNC should be different from the case where the failed IP link is a single-layer IP link under P-PNC responsibility.

Comment: The assumption in this text is that there are grey interfaces between the routers and the optical NEs. More investigation is needed for the scenarios where optical pluggable interfaces are used in the router. Three scenarios for WDM networks: grey interfaces, colored interfaces option 1 and colored interfaces option 2. To consider also the case with ODU switching.

5.3. Cross-layer Link Failures

The failures discussed in this section occur on the connection between a router and a ROADM. A first case concerns the Tx fiber used by R1 to send traffic to ROADM1 (Figure 4).

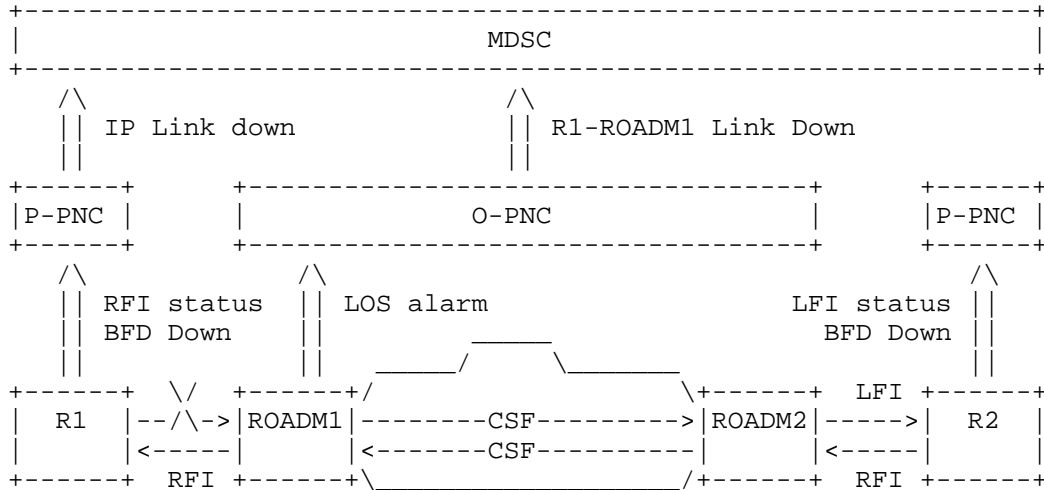


Figure 4: Failure on the optical ingress link

The failure on that fiber is physically detected by ROADM1 that sends a corresponding notification to O-PNC and generates a Client Signal Fail (CSF) message along the optical path. Upon receiving CSF, ROADM2 sends a LFI to R2. If R2 is instructed to decode physical transmission messages, upon receiving LFI it generates a corresponding message to P-PNC, also informing of the loss of IP connectivity due to missed Bidirectional Forwarding Detection (BFD) messages [RFC5880]. At the physical level, R2 may generate on its Tx interface an RFI indication that is propagated downstream to the optical network (where a CSF is generated) and to R1. When receiving RFI on its Rx interface, R1 can also send a notification to P-PNC. When O-PNC and P-PNC get the notifications sent by the network elements, they also instruct MDSC. O-PNC informs MDSC of the link R1-ROADM1 down, while P-PNC informs MDSC that the corresponding IP link is down due to missed BFD signalling in addition to the RFI status. It is up to the MDSC to correlate the events and determine what IP services are affected (VPNs, P2P links, etc.).

A second case is depicted in figure Figure 5. The failure happens on the Rx fiber used by R2 to receive traffic from ROADM2.

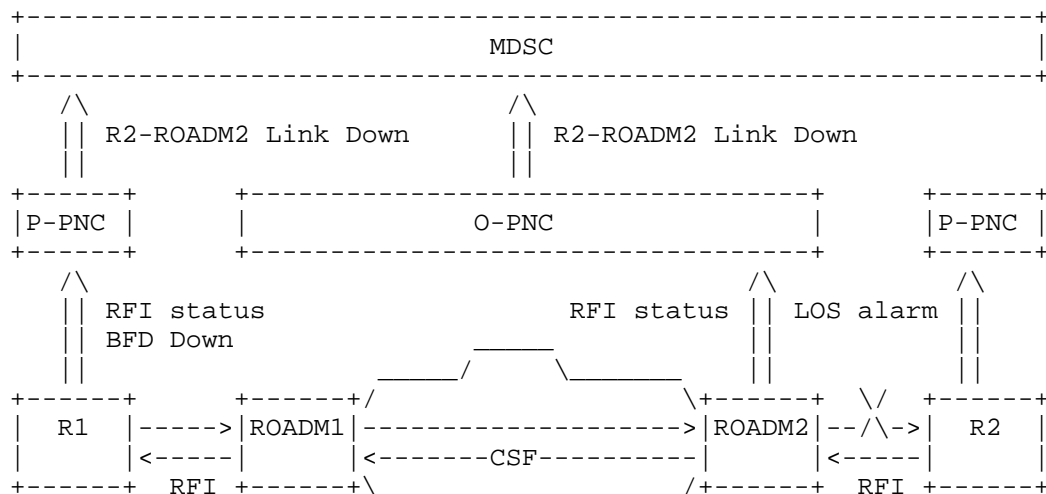


Figure 5: Failure on the optical egress link

R2 physically detects the absence of signal generating a corresponding LOS alarm to P-PNC. In turn, P-PNC signals MDSC of the corresponding event affecting the link between ROADM2 and R2. R2 also propagates an RFI indication on the return fiber. Upon detecting it, ROADM2 also informs O-PNC of the failure with a corresponding RFI status indication. ROADM2 also propagates a CSF indication across the optical domain, translated to an RFI by ROADM1 towards R1. The RFI is detected by R1 that may inform P-PNC about the remote failure with an RFI status indication, if instructed to do so, and with a BFD down event notification when detecting missing connectivity. As noted, MDSC correlates the events to determine the affected services.

A failure may also occur when the two unidirectional fibers connecting a router, e.g. R1, to a ROADM, e.g. ROADM2, are affected, for example for a simultaneous fiber cut, as shown in figure Figure 6.

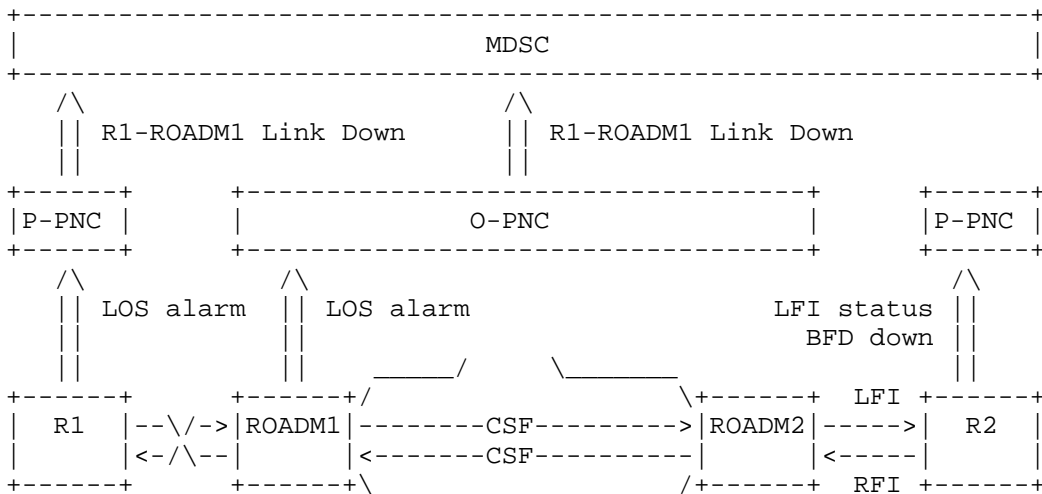


Figure 6: Failure on the access link

Both Tx and Rx fibers are affected, then R1 and ROADM1 immediately detect physical LOS and inform P-PNC and O-PNC respectively. ROADM1 also triggers a CSF indication towards the optical core that eventually gets to ROADM2, which sends a LFI to R2. R2 may detect this signal, informing P-PNC. From the IP connectivity standpoint, based on the configured BFD detection parameters, R2 also signals to P-PNC the lack of end-to-end connectivity. It then generates a RFI indication back to ROADM2, which in turn sends a CSF indication on the optical return path. Both P-PNC and O-PNC inform MDSC of the event affecting the link between R1 and ROADM1 for its successive correlation.

5.4. Router Node Failures

In this case it is assumed that a router port experiences a hardware failure, for example R1's port connecting to ROADM1. R1 may have internal mechanisms that detect the failure and trigger the relevant notification to P-PNC. At the IP level, the missing reception of the BFD messages against R2 triggers a BFD down notification to P-PNC. The same notification is sent by R2, confirming that the IP connectivity is lost.

5.5. VPN Service Fault Monitoring

The fault detection and localization mechanisms described in the preceding subsections feed into a broader process in which the MDSC takes remedial action to restore affected VPN services and, once resolution is achieved, notifies the Orchestrator layer.

When a failure occurs, it is detected by a network element (router, ROADM, or a node within the optical core) and the corresponding notification (alarm, BFD down, LOS, RFI, etc.) is sent to the responsible PNC (P-PNC or O-PNC), which in turn notifies the MDSC via the MPI. The MDSC then correlates the incoming notifications from the P-PNC(s) and O-PNC, determines the root cause and scope of the failure, and identifies which VPN services (L2VPN/L3VPN) are affected. The subsequent remedial actions are described in Section 5.5.1.

5.5.1. Remedial Action and Resolution

Based on the failure type and layer, the MDSC requests new connectivity from the appropriate PNC or combination of PNCs:

- * If the failure is confined to the IP layer (e.g. a router port failure), the MDSC requests the P-PNC to reroute the affected IP traffic to an alternative IP path.
- * If the failure is at the optical layer (e.g. a fiber cut or ROADM failure), the MDSC requests the O-PNC to establish a new optical path (via restoration or protection mechanisms) and subsequently instructs the P-PNC to steer IP traffic over the newly available optical connection.
- * If the failure spans both layers, the MDSC coordinates sequential actions on both the O-PNC (new optical path) and the P-PNC (IP traffic switch), ensuring consistency across layers.

Once connectivity is restored and the VPN services are operational again, the P-PNC and/or O-PNC notify the MDSC of the successful recovery. The MDSC then informs the Orchestrator layer of the updated service status and the resolution of the fault. The specific mechanisms used for this notification between the MDSC and the Orchestrator are outside the scope of this document.

The detailed step-by-step sequences for each failure scenario and the corresponding multi-layer recovery procedures are described in Section 7.

6. Multi-layer Performance Management

Network performance management refers to the set of operational actions taken to resolve issues affecting network performance that may degrade the quality of services offered to customers.

For the scope of the present document, which focuses on multi-layer, multi-domain networks, two cases are of interest: 1. The optical layer detects, through performance data measurement, collection, and analysis, that an abnormal condition (e.g., physical signal degradation) has arisen or is likely to happen in either of the optical domains considered in Figure 1. The O-PNC provides relevant information to the MDSC (e.g., the fiber where the degradation was detected), which triggers correlation analysis by the MDSC to detect whether any services are impacted at the IP level and, if so, to take corrective actions through the P-PNC (e.g., traffic rerouting). 2. The IP layer detects, through performance data measurement, collection, and analysis, that the SLA associated with transport of a VPN service is not compliant in at least one of the two IP domains represented in Figure 1. The P-PNC provides relevant information to the MDSC (e.g., the IP tunnel carrying the VPN service), which enables the MDSC to take reactive measures with support from the P-PNC (e.g., rerouting IP traffic onto a different IP path). The MDSC can take further steps, such as verifying through the O-PNC whether any failure or degradation has happened in the optical layer, but this is out of scope for case 2. The focus here is on IP multi-domain, end-to-end performance management.

The two cases are further detailed in the relevant subsections.

6.1. Optical performance management

Optical devices employ mechanisms for monitoring the condition of an OTN link. Among others, pre-Forward Error Correction (pre-FEC) Bit Error Rate (BER) allows tracking bit errors on the optical wire, notifying the transmitter side or a controlling agent when a specified threshold is reached or passed. The advantage of this mechanism is to get an early warning on the optical path performance: exceeding the specified threshold means that the receiver is no longer able to correct all the errors on the channel. As a result, the transmitter or the controlling entity (e.g. an SDN controller) may trigger counter-actions such as the switch to a different optical path.

In the context of multi-layer performance management, it is assumed that: 1. The O-PNC is capable of monitoring the optical performance of DWDM links and generating a TCA toward the MDSC when the pre-FEC BER value exceeds a user-specified threshold. 2. The MDSC is capable of correlating the pre-FEC BER TCA with a related IP link and taking appropriate corrective actions, if programmed to do so.

In this context, the assumption is that pre-FEC BER measurement is done on the optical path between ROADM1 and ROADM2 of Figure 2. Some IP services (e.g. L2/L3 VPNs) are active between R1 and R2, using

the optical path between ROADM1 and ROADM2 as a transport. The sequence of steps to handle the exception detected by the optical performance management is expected to be the following:

- * step 1. ROADM2 detects a pre-FEC BER value at an ingress interface higher than the defined threshold. A corresponding alarm is sent to the O-PNC
- * step 2. O-PNC forwards the alarm to MDSC
- * step 3. MDSC correlates the information of the optical path subject to pre-FEC BER issues and the IP services active on it.

Depending on how the MDSC is instructed to react, different choices are possible. At one extreme of the spectrum, the MDSC notes the event and simply triggers a notification to the operator. At the other extreme, the MDSC may start the multi-layer resiliency mechanisms described in Section 5.2, as the case is equivalent to the handling of an optical failure.

6.2. End-to-end IP performance management

Performance measurement at the IP layer may be based on a multiplicity of methods, including interface counters, passive and active mechanisms [RFC7799]. While the utilization of those mechanisms is not constrained by network topology, for example by the number of IP domains crossed by a measurement flow, in practice they are often enabled in limited environments (controlled domains) [RFC8799].

As a result, the applicability of such methods is often limited to a single IP domain due to the necessity of avoiding the exchange and disclosure of sensitive data across multiple administrative organizations. With reference to Figure 1, it is then assumed that both IP domains, namely Packet domain 1 and 2, run separate performance measurements. It is the responsibility of each P-PNC to inform the MDSC in the case of service SLA degradation so that the MDSC enables a corrective action.

7. Multi-layer Resiliency

The coordination of both the IP and the optical layer in the cases discussed in Section 7 requires the MDSC to be aware of some network capabilities and to exchange the corresponding information with both the P-PNC and the O-PNC.

To achieve maximum flexibility, a network operator may enable or disable these capabilities. Once the network operator has configured the capabilities described in this section, the MDSC exchanges the relevant configuration with the PNCs present in the network before the use cases described in Section 7 take place.

The list of parameters that the MDSC may need to communicate to the PNCs includes:

- * IP service reversion: on/off
- * Optical service reversion: on/off
- * Hold-off time: time in ms (0 for immediate fast re-routing)
- * Wait time before reversion: time in s
- * Recovery method used in the optical layer: protection/restoration

7.1. Optical Network Failures

Failures in the optical domain can be recovered by packet-based protection mechanisms as described in [I-D.ietf-teas-actn-poi-applicability].

This use case is characterized by a fault happening on the upper fiber connecting ROADM1 and ROADM2 (port P3 to port P3 as depicted in Figure 2), affecting the IP traffic between R1 and R2. As a result, the MDSC and the domain controllers cooperate to find a backup path for the IP traffic. If the optical layer does not employ any mechanisms, the case is typically solved through the Fast Rerouting Mechanisms (FRR) enabled by the IP/MPLS control plane. With reference to figure Figure 2, this corresponds to using the combination of the two detour paths R1-R3 and R3-R2. For the scope of this document, the assumption is instead that the optical layer supports its own mechanisms that have to interact with the IP layer. Two sub-cases are possible:

1. The optical layer supports restoration
2. The optical layer supports protection.

7.1.1.1. Optical restoration

As restoration typically sets an alternative path on the fly based on the availability of sufficient optical resources, the time taken by the process to create an optical backup tends to be longer than the time taken by the IP/MPLS FRR process. As a result, the interaction between the two layers follows the mimics shown in the next figure.

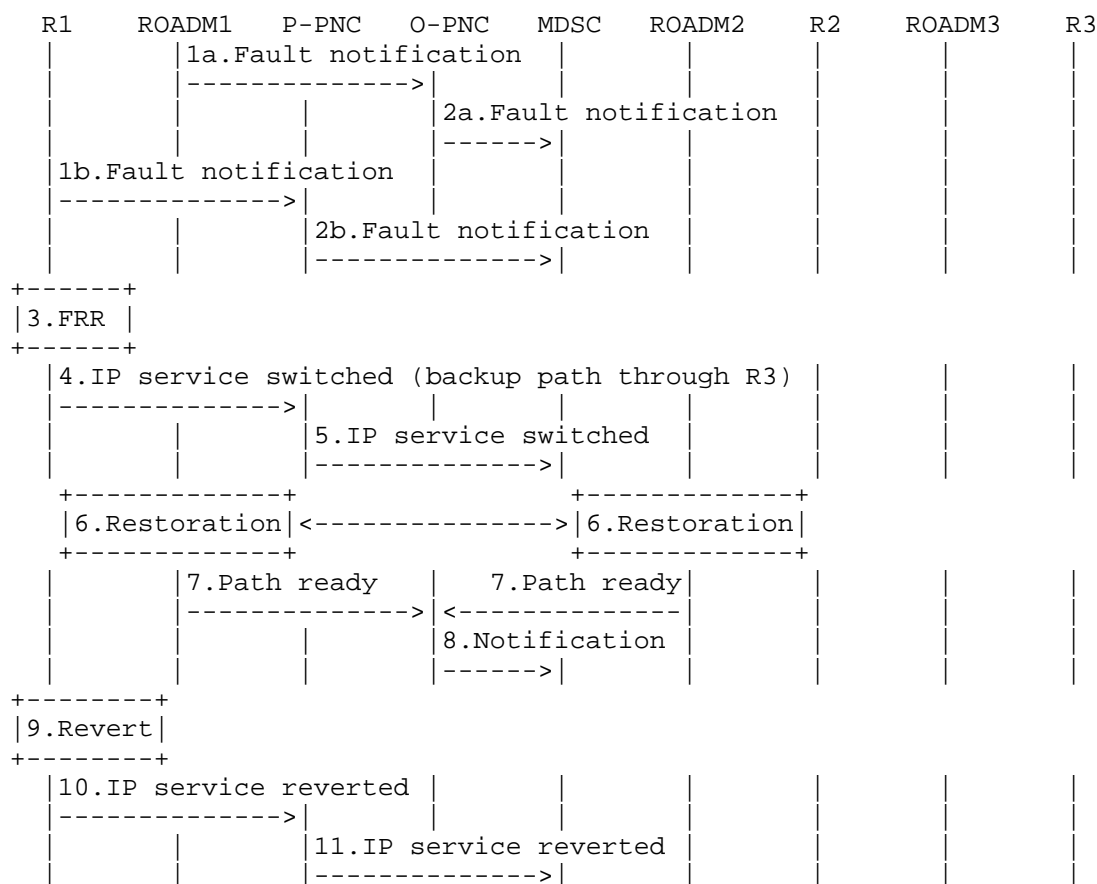


Figure 7: Fault detection with optical restoration

More in details:

- * step 1a. The fault on the optical path (e.g. fiber cut, loss of signal, etc.) is detected by ROADM1 and notified to O-PNC
- * step 2a. O-PNC notifies the fault to MDSC

- * step 1b. R1 detects loss of end-to-end connectivity (e.g. 3 missed BFD messages) and notifies P-PNC. This step takes place almost simultaneously to 1a.
- * step 2b. P-PNC notifies the issue to MDSC
- * step 3. R1 starts a fast reroute process to enable a backup path at the IP/MPLS layer, using the already established detour through R3
- * step 4. R1 notifies P-PNC of the IP service switch through the alternate path (R1-R3 and R3-R2)
- * step 5. P-PNC notifies MDSC of the switch
- * step 6. ROADM1 and ROADM2 enable the restoration process. Based on the mechanism adopted, there may be interaction between them
- * step 7. Both ROADM1 and ROADM2 notify O-PNC of the availability of an optical backup path
- * step 8. O-PNC notifies MDSC of the availability of an optical backup path
- * step 9. R1 detects again end-to-end connectivity through the initial path R1-R2 and, if configured to do so, revert the service
- * step 10. R1 notifies P-PNC of the switch to the initial path
- * step 11. P-PNC notifies the switch to MDSC.

As noted in step 6., the restoration process may require an exchange of messages between ROADM1 and ROADM2. This is not detailed in the present document as it is assumed that the relevant signaling is handled through O-PNC.

In step 9., R1 detects again control traffic from R2. The decision whether to revert the service on the initial path is local, e.g. it depends on the configuration made by the network operator. Often, the IP equipment is configured to operate the reversion automatically, but there are cases where the network operator may prefer differently.

At the end of the process, multi-layer hitless reversion may take place, again based on the configuration adopted by the network operator. If multi-layer hitless reversion is adopted, then the process described in Section 7.5 takes place.

7.1.2. Optical protection

Differently from the previous case, here optical protection is considered. The duration of this process is comparable with IP/MPLS FRR, as it is pre-computed. As a consequence, when multi-layer coordination is enabled it is preferable to hold-off FRR on R1 and wait that optical protection is completed. The process is shown in the next figure.

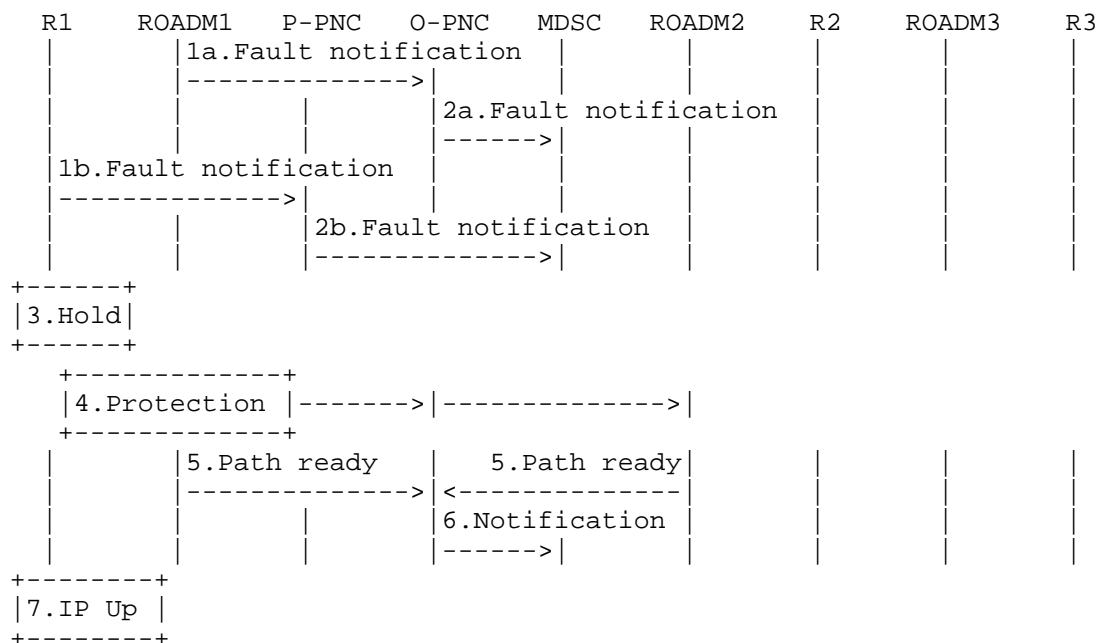


Figure 8: Fault detection with optical protection

The detailed process includes the following steps:

- * step 1a. The fault on the optical path (e.g. fiber cut, loss of signal, etc.) is detected by ROADM1 and notified to O-PNC
- * step 2a. O-PNC notifies the fault to MDSC
- * step 1b. R1 detects loss of end-to-end connectivity (e.g. 3 missed BFD messages) and notifies P-PNC. This step takes place almost simultaneously to 1a.
- * step 2b. P-PNC notifies the issue to MDSC [Editor's note: is this step necessary?]

- * step 3. R1 is configured to hold the FRR process, thus it waits for the corresponding value set by the hold-off time parameter
- * step 4. Optical protection is started by ROADM1, potentially involving an exchange of messages with O-PNC and ROADM2
- * step 5. Both ROADM1 and ROADM2 notify O-PNC of the availability of an optical backup path
- * step 6. O-PNC notifies MDSC of the availability of an optical backup path
- * step 7. R1 detects again end-to-end connectivity with R2.

The IP traffic is recovered as soon as the optical protection is completed with no action taken by the IP routers.

As in the previous use case, when the failure is fixed the network operator may desire to bring the service back to the original configuration. If this is the case, multi-layer hitless reversion, as described in Section 7.5, takes place to move the service back to the initial network setup.

7.2. Optical Network Maintenance

Before planned maintenance operation on the optical network takes place, the IP traffic affected by the maintenance operation should be moved hitlessly to another link. The MDSC and the P-PNC have to coordinate to reroute the traffic before the event happens. In such a case the IP traffic needs to be locked to the protection route until the maintenance event is finished, unless a fault occurs on such path. In this example, it is supposed that the link undergoing maintenance activity is the one from ROADM1 to ROADM2, affecting the IP traffic steered from R1 to R2. A few minutes before the maintenance window, the MDSC starts the process that brings to the hitless re-routing of the affected IP traffic. That means the IP backup path (through R3) is available and it is used only for the time requested by the optical plane to do maintenance. The path R1-R3 should not be overloaded, unless the network operator accepts some possible traffic losses. At the optical layer, maintenance has no traffic impact because a new path is configured in advance, and the optical service does not revert to the original link until the maintenance window is finished. At the end of maintenance, the network configuration is moved back to the initial setup, using the multi-layer hitless reversion process discussed in Section 7.5 if the network operator has enabled it.

- * step 2. P-PNC signals R1 to switch IP service to the backup path
- * step 3. R1 switches to backup path and acks to P-PNC
- * step 4. P-PNC acks to MDSC
- * step 5. MDSC instructs O-PNC to enable the process to create an optical backup path
- * step 6. O-PNC instructs ROADM1 and ROADM2 to enable a backup path
- * step 7. ROADM1 and ROADM2 acknowledge to O-PNC
- * step 8. O-PNC acknowledges to MDSC
- * step 9. MDSC instructs O-PNC to disable the primary optical path, initially used, and switch to the optical backup path
- * step 10. O-PNC instructs ROADM1 and ROADM2 to switch
- * step 11. ROADM1 and ROADM2 acknowledge to O-PNC
- * step 12. O-PNC acknowledges to MDSC
- * step 13. MDSC requires P-PNC to revert the IP service back to the primary path (R1-R2)
- * step 14. P-PNC signals R1 to switch IP service to the primary path (carried over the optical backup path)
- * step 15. R1 switches to the backup path and acknowledges to P-PNC
- * step 16. P-PNC acknowledges to MDSC
- * step 17. The maintenance activity follows.

Once the activity is over, the network operator may wish to bring the whole configuration back to the IP and optical primary paths. In such a case, multi-layer hitless reversion may be performed, as described in Section 7.5.

7.3. Cross-layer Link Failures

The approach described here leverages multi-layer POI capabilities to address failures in links between IP routers and ROADMs, relying on optical network protection/restoration to handle most failure scenarios. The connectivity between a router and an edge ROADM is characterized by having N working ports and one spare port (N+1) for protection.

Depending on the specific network configuration and protection scheme adopted, this approach may offer cost advantages because it reduces the overall resources required for protection in the optical network. Since the number of failed links between IP routers and edge ROADMs is lower, this configuration can achieve higher availability at lower cost while recovering 100% of IP traffic.

Following the previous examples, this case is characterized by R1 being configured with N working ports (say, P1-P3) and one spare port (PP) as protection for the other N ports. In case of failure, for example on port P1, PP is dynamically activated and traffic originally directed to P1 is steered to PP. PP receives the same configuration as P1, while P1 is brought to a down state. Unlike ordinary LAG, traffic is not redistributed over the surviving links. Since a backup port (PP) is enabled, traffic continues to flow on N links instead of N-1. Although this introduces additional IP-layer complexity by requiring an extra port on both R1 and ROADM1, the optical-layer configuration shown in figure Figure 2 does not change, because only N optical channels (e.g. lambdas) are used, as shown in figure Figure 10.

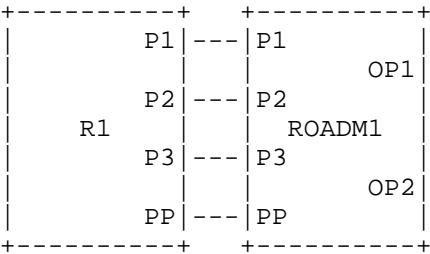


Figure 10: Use of N:1 protection on R1

Two sub-cases may be considered, depending on the availability of a Muxponder or a Transponder on ROADM1. If a Muxponder is used, then the optical P1 and PP are hosted on the same optical complex (e.g. board) on the customer's edge of ROADM1. It is the optical complex that selects the input source of the signals and maps it on the proper lambda. If instead a Transponder is used, then it's ROADM1's

internal matrix that switches from the input source from P1 to PP, cross-connecting the signal to the output lambda. It has to be noted that the mechanism to deal with the on-the-fly reconfiguration of a router's port is out of the scope of the present document and may be subject of a dedicated draft.

The next figure shows the process adopted to handle N:1 port protection.

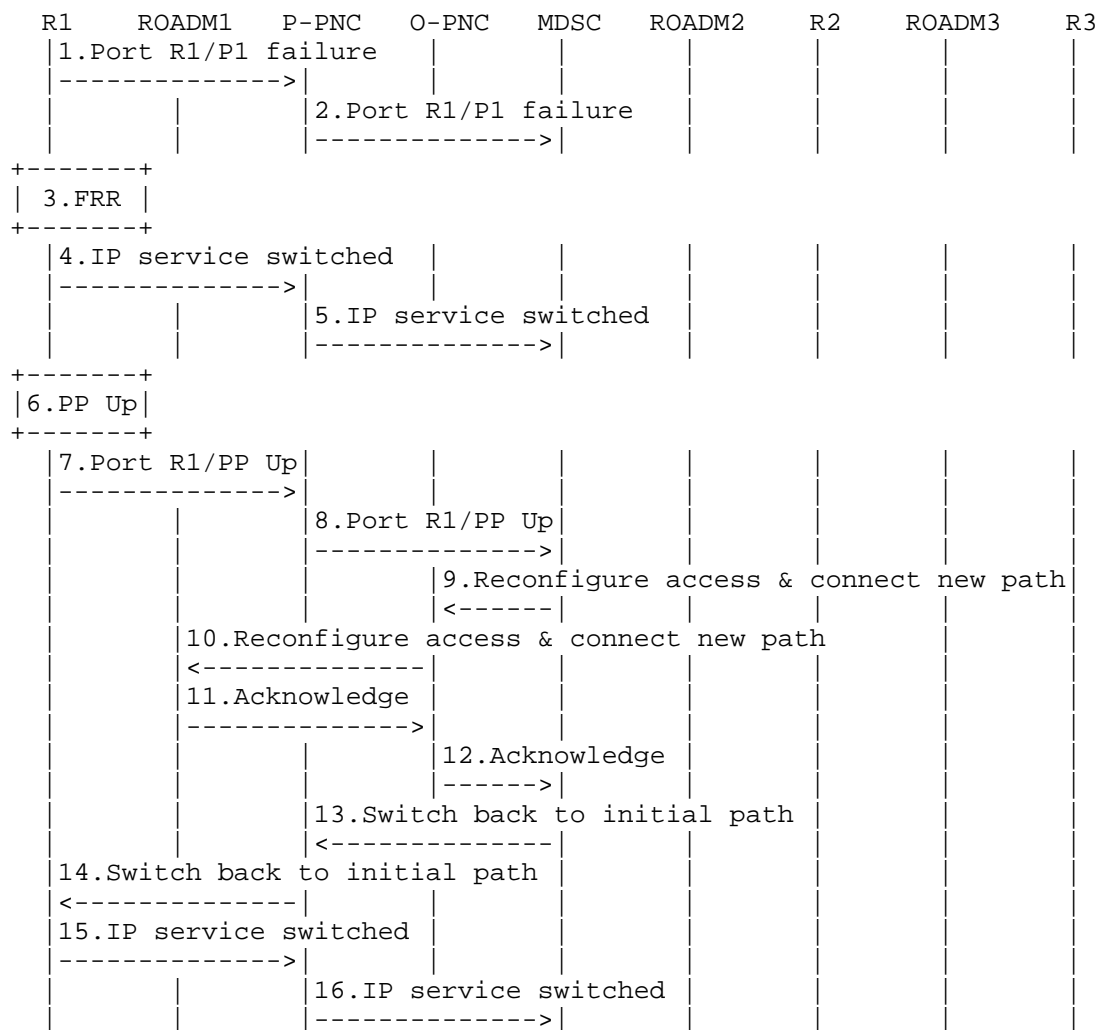


Figure 11: N:1 protection operation

The sequence of steps is detailed.

- * step 1. R1 detects port P1 failure and notifies P-PNC
- * step 2. P-PNC notifies MDSC of the failure
- * step 3. R1 triggers FRR to protect the IP flows steering
- * step 4. R1 informs P-PNC of the switch to the backup path
- * step 5. P-PNC notifies MDSC of the traffic switch
- * step 6. R1 handles the mechanism to replicate the configuration of P1 to PP
- * step 7. R1 informs P-PNC that PP is up and ready to forward traffic
- * step 8. P-PNC notifies MDSC that port PP is up and ready to forward traffic
- * step 9. MDSC requires O-PNC to reconfigure ROADM1 access (both in the case of muxponder and transponder) and WDM connectivity if a transponder is used
- * step 10. O-PNC signals ROADM1 to reconfigure access (muxponder/transponder) and WDM connectivity (transponder)
- * step 11. ROADM1 acknowledges to O-PNC
- * step 12. O-PNC acknowledges to MDSC
- * step 13. MDSC requires P-PNC to revert to the initial (primary) path
- * step 14. P-PNC notifies R1 to revert to initial (primary) path
- * step 15. R1 notifies P-PNC of IP service switch and new port in use
- * step 16. P-PNC notifies MDSC of service switch and new port in use

As in the previous cases, when port P1 on R1 is fixed, multilayer reversion Section 7.5 to the initial configuration may happen. This is dependent on the network operator's preference.

7.4. Router Node Failures

As shown in Figure 2, in its normal operations R1 is dual-homed to R2 and R3. Even if highly unlikely due to the usual redundancy deployed in field, this case considers a full failure of R2 (node failure). The implications of such an event are useful to discuss the interaction between the IP and the optical layers through the MDSC coordination. The underlying assumption is that it is not possible to R2 to communicate to P-PNC about the event causing the failure, so it is up to R1 to detect it and to communicate instead to P-PNC. The first reaction to the event is to perform a fast-rerouting action and move the traffic from the R1-R2 link to the R1-R3 link. As part of the assumption, the R1-R3 IP link has been previously dimensioned to carry a certain amount of traffic, so it is possible that after fast re-routing takes place some traffic previously carried on the R1-R2 IP link and now shifted to R1-R3 is discarded, for example because congestion occurs. MDSC instructs the optical layer to find available optical resources, activate a new optical path between ROADM1 and ROADM3 and finally move the traffic previously associated to R1-R2 to the newly created optical path. When this second optical path is available, MDSC triggers a new switch of the traffic so that R1 can now steer the previous R1-R2 traffic to the new optical path. The final configuration is shown in figure Figure 12.

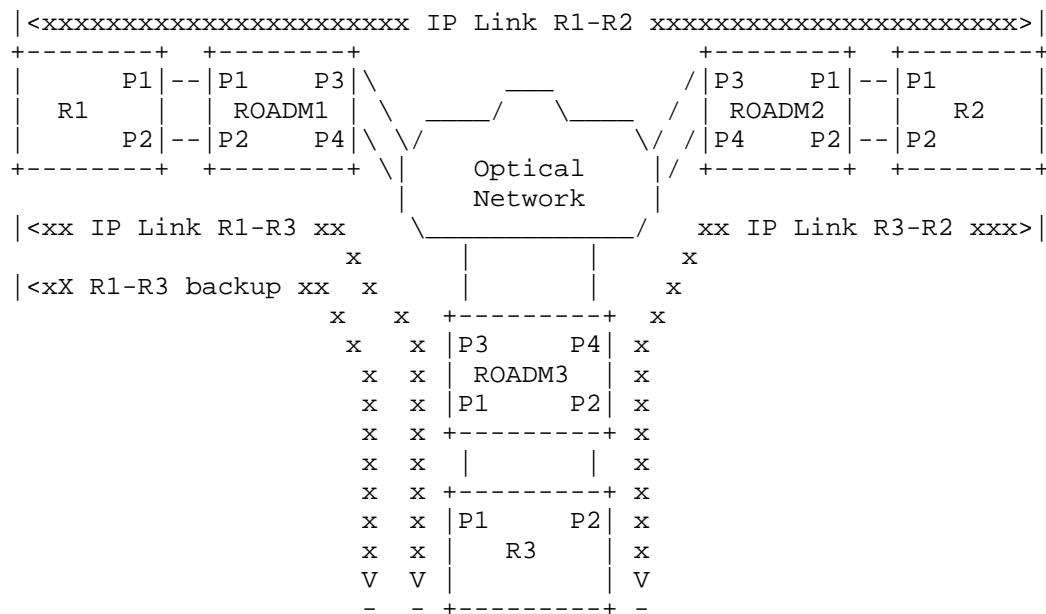


Figure 12: IP configuration after the creation of a second optical path

The next figure shows the process adopted to handle the node protection case.

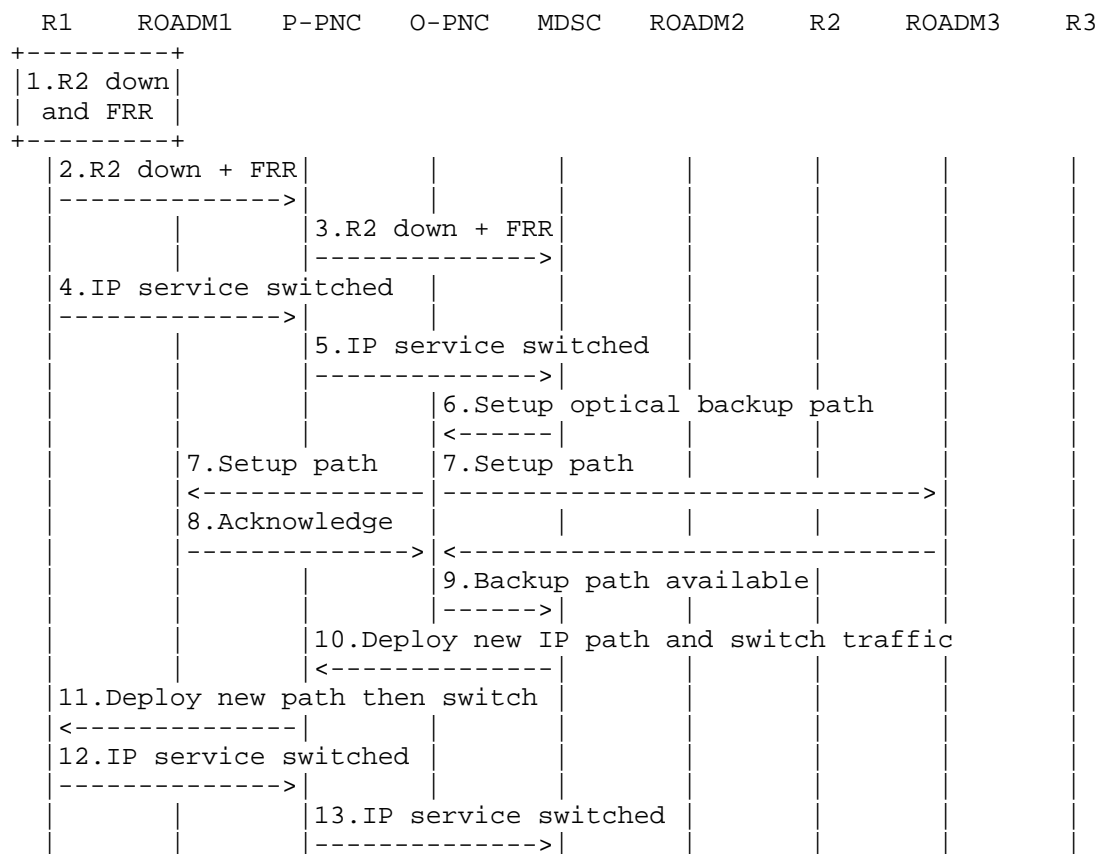


Figure 13: Node protection operation

- * step 1. R1 detects R2's failure and triggers IP FRR finding R3 as the next hop
- * step 2. R1 notifies P-PNC that R2 is down and FRR has started
- * step 3. P-PNC notifies MDSC of the events
- * step 4. Upon moving the R1-R2 traffic (or part of it) on R1-R3 path, R1 notifies P-PNC of the service switch
- * step 5. P-PNC notifies MDSC of the switch

- * step 6. MDSC requires O-PNC to compute a new optical path between ROADM1 and ROADM3
- * step 7. O-PNC instructs both ROADM1 and ROADM3 to configure a new optical service
- * step 8. Both ROADM1 and ROADM3 inform O-PNC that the backup path is available
- * step 9. O-PNC informs MDSC that the backup path is available
- * step 10. MDSC computes a new IP path between R1 and R3, provides the relevant information to P-PNC and triggers switch
- * step 11. P-PNC transfers the information received to R1 and triggers R1 to switch traffic
- * step 12. R1 informs P-PNC of the service switch
- * step 13. P-PNC informs MDSC of the service switch.

7.5. Multi-layer hitless reversion

In some cases, the mechanisms employed by the optical layer to revert to the original setup may cause disruption at the IP layer, if proper coordination is not enabled. As this may cause traffic loss, if the optical reversion is requested by the network operator, multi-layer coordination under the supervision of the MDSC is necessary. The effect of multi-layer coordination is to bring the whole network, i.e. both the IP and the optical layers, back to their initial configuration after the recovery from a failure. In particular, the process described in this section relies on the hitless switching capability of the IP layer. Depending on the specific configuration, the procedure can be enabled at the end of the use cases described in Section 7. The decision whether to apply it or not has to be evaluated by the network operator considering different factors, including the relative complexity of the process and the effects of its steps on the live traffic.

To move back to the initial network configuration the MDSC has to follow a sequence of steps:

- * Force the IP layer to switch the traffic flow(s) on another path, e.g. an alternative/backup path
- * Trigger the optical layer to coordinate the reversion to the initial setup, e.g. disable an optical backup path and enable connectivity on the previously used primary path

- * Force again the IP layer to switch back to the original path. The actions on the IP layer are handled so that the IP traffic is switched only after the interface queues are emptied, guaranteeing a hitless switching.

The mimics of the steps requested is shown in the next figure.

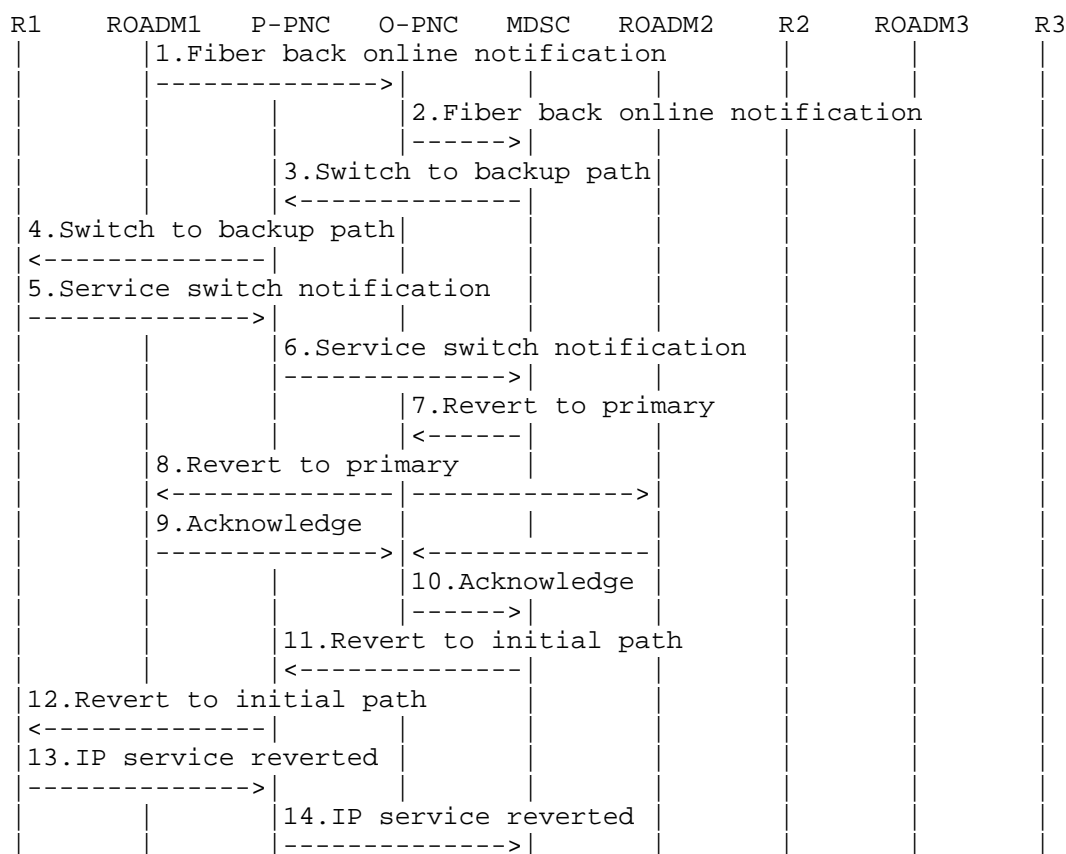


Figure 14: hitless multi-layer reversion

The steps illustrated in the previous figure are detailed here:

- * step 1. ROADM1 detects the optical signal is up again on the previously broken fiber and notifies O-PNC
- * step 2. O-PNC notifies MDSC of the fiber up event

- * step 3. MDSC requires P-PNC to move the affected IP service(s) to an alternative/backup path (this path may vary according to the scenarios explained later). Being a hitless switch, it is necessary to avoid loss of service
- * step 4. P-PNC signals R1 to switch the IP service(s) to the alternative/backup path
- * step 5. R1 switches the service(s) to the alternative/backup path and notifies P-PNC
- * step 6. P-PNC confirms the switch to MDSC
- * step 7. MDSC instructs O-PNC to disable the optical protection path (which may vary according to the scenarios detailed later) and activate again the optical primary path
- * step 8. O-PNC instructs both ROADM1 and ROADM2 to disable the optical protection path and activate the primary one
- * step 9. ROADM1 and ROADM2 acknowledge to O-PNC
- * step 10. O-PNC acknowledges to MDSC
- * step 11. MDSC requires P-PNC to revert the IP service(s) back to the primary path
- * step 12. P-PNC signals R1 to switch the IP service(s) to primary path
- * step 13. R1 switches and acknowledges to P-PNC
- * step 14. P-PNC acknowledges to MDSC.

8. Conclusions

This section will provide a summary of the analysis and of the gaps identified in this draft once the analysis is mature.

9. Security Considerations

This document analyses service assurance procedures across multiple ACTN control entities and management interfaces. Compromise or misuse of these interfaces can lead to incorrect correlation of faults, unauthorized remedial actions, or service disruption across packet and optical layers.

Implementations and deployments SHOULD ensure authentication, authorization, and integrity protection for exchanges across the MPI and related controller interfaces. Access to assurance telemetry, alarms, and topology-correlated state SHOULD be limited to authorized roles, and sensitive operational data SHOULD be protected in transit and when stored.

Because this document includes multi-domain scenarios, operators should evaluate trust boundaries between domains and apply policy controls to reduce risks from spoofed alarms, stale telemetry, and denial-of-service conditions affecting assurance workflows.

10. IANA Considerations

This document has no IANA actions.

11. References

11.1. Normative References

[I-D.feng-opsawg-incident-management]

Feng, C., Hu, T., Contreras, L. M., Graf, T., Wu, Q., Yu, C., and N. Davis, "Incident Management for Network Services", Work in Progress, Internet-Draft, draft-feng-opsawg-incident-management-04, 30 January 2024, <<https://datatracker.ietf.org/doc/html/draft-feng-opsawg-incident-management-04>>.

[I-D.ietf-teas-actn-poi-applicability]

Peruzzini, F., Bouquier, J., Busi, I., King, D., and D. Ceccarelli, "Applicability of Abstraction and Control of Traffic Engineered Networks (ACTN) to Packet Optical Integration (POI)", Work in Progress, Internet-Draft, draft-ietf-teas-actn-poi-applicability-17, 22 December 2025, <<https://datatracker.ietf.org/doc/html/draft-ietf-teas-actn-poi-applicability-17>>.

[I-D.yu-performance-monitoring-yang]

Yu, C., "A YANG Data Model for Optical Performance Monitoring", Work in Progress, Internet-Draft, draft-yu-performance-monitoring-yang-00, 24 October 2022, <<https://datatracker.ietf.org/doc/html/draft-yu-performance-monitoring-yang-00>>.

[ITU-T_G.709]

International Telecommunication Union, "Interfaces for the optical transport network", ITU-T Recommendation G.709, Amendment 3 , March 2024,
<<https://www.itu.int/rec/T-REC-G.709/>>.

[ITU-T_G.7710]

International Telecommunication Union, "Common equipment management function requirements", ITU-T Recommendation G.7710, Amendment 1 , November 2022,
<<https://www.itu.int/rec/T-REC-G.7710/>>.

[ITU-T_G.798]

International Telecommunication Union, "Characteristics of optical transport network hierarchy equipment functional blocks", ITU-T Recommendation G.798 , April 2024,
<<https://www.itu.int/rec/T-REC-G.798/>>.

[ITU-T_G.874]

International Telecommunication Union, "Management aspects of optical transport network elements", ITU-T Recommendation G.874, Amendment 2 , January 2024,
<<https://www.itu.int/rec/T-REC-G.874/>>.

[RFC8453]

Ceccarelli, D., Ed. and Y. Lee, Ed., "Framework for Abstraction and Control of TE Networks (ACTN)", RFC 8453, DOI 10.17487/RFC8453, August 2018,
<<https://www.rfc-editor.org/rfc/rfc8453>>.

[RFC8632]

Vallin, S. and M. Bjorklund, "A YANG Data Model for Alarm Management", RFC 8632, DOI 10.17487/RFC8632, September 2019, <<https://www.rfc-editor.org/rfc/rfc8632>>.

[RFC9182]

Barguil, S., Gonzalez de Dios, O., Ed., Boucadair, M., Ed., Munoz, L., and A. Aguado, "A YANG Network Data Model for Layer 3 VPNs", RFC 9182, DOI 10.17487/RFC9182, February 2022, <<https://www.rfc-editor.org/rfc/rfc9182>>.

[RFC9291]

Boucadair, M., Ed., Gonzalez de Dios, O., Ed., Barguil, S., and L. Munoz, "A YANG Network Data Model for Layer 2 VPNs", RFC 9291, DOI 10.17487/RFC9291, September 2022, <<https://www.rfc-editor.org/rfc/rfc9291>>.

[RFC9375]

Wu, B., Ed., Wu, Q., Ed., Boucadair, M., Ed., Gonzalez de Dios, O., and B. Wen, "A YANG Data Model for Network and VPN Service Performance Monitoring", RFC 9375, DOI 10.17487/RFC9375, April 2023,
<<https://www.rfc-editor.org/rfc/rfc9375>>.

- [RFC9418] Claise, B., Quilbeuf, J., Lucente, P., Fasano, P., and T. Arumugam, "A YANG Data Model for Service Assurance", RFC 9418, DOI 10.17487/RFC9418, July 2023, <<https://www.rfc-editor.org/rfc/rfc9418>>.

11.2. Informative References

- [I-D.ietf-ccamp-actn-optical-transport-mgmt]
Tan, XingZhao, Yu, C., King, D., and A. Farrel,
"Integrating YANG Configuration and Management into an
Abstraction and Control of TE Networks (ACTN) System for
Optical Networks", Work in Progress, Internet-Draft,
draft-ietf-ccamp-actn-optical-transport-mgmt-04, 30
November 2025, <[https://datatracker.ietf.org/doc/html/
draft-ietf-ccamp-actn-optical-transport-mgmt-04](https://datatracker.ietf.org/doc/html/draft-ietf-ccamp-actn-optical-transport-mgmt-04)>.
- [I-D.mix-teas-actn-poi-extension]
Galimberti, G., Bouquier, J., Gerstel, O., Foster, B., and
D. Ceccarelli, "Applicability of Abstraction and Control
of Traffic Engineered Networks (ACTN) to Packet Optical
Integration (POI) extensions to support Router Optical
interfaces.", Work in Progress, Internet-Draft, draft-mix-
teas-actn-poi-extension-00, 24 October 2022,
<[https://datatracker.ietf.org/doc/html/draft-mix-teas-
actn-poi-extension-00](https://datatracker.ietf.org/doc/html/draft-mix-teas-actn-poi-extension-00)>.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection
(BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010,
<<https://www.rfc-editor.org/rfc/rfc5880>>.
- [RFC7799] Morton, A., "Active and Passive Metrics and Methods (with
Hybrid Types In-Between)", RFC 7799, DOI 10.17487/RFC7799,
May 2016, <<https://www.rfc-editor.org/rfc/rfc7799>>.
- [RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet
Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020,
<<https://www.rfc-editor.org/rfc/rfc8799>>.

Acknowledgments

TODO acknowledge.

Contributors

Daniel King
Old Dog Consulting
Email: daniel@olddog.co.uk

Authors' Addresses

Italo Busi
Huawei Technologies
Email: italo.busi@huawei.com

Jean-Francois Bouquier
Vodafone
Email: jeff.bouquier@vodafone.com

Fabio Peruzzini
FiberCop
Email: fabio.peruzzini@fibercop.com

Paolo Volpato
Huawei Technologies
Email: paolo.volpato@huawei.com

Prasenjith Manna
Cisco
Email: prmanna@cisco.com