

INTERNET-DRAFT  
Intended Status: Informational  
Expires: May 5, 2013

G. Mattson  
T. Nadeau  
Juniper Networks  
November 5, 2012

**IRS Use Case for IP and Transport Workflows**  
**draft-mattson-irs-usecase-00**

Abstract

This document describes use-cases for IRS to provide a lightweight method for common management and control state for typical operations and workflows of multilayer networks involving both IP and sub-IP transport. IRS provides a lightweight, streaming interface between routing systems and external applications. By extending the IRS model to transport systems, multi-layer networks can be managed and used more efficiently. IRS may enable the integration of IP and transport maintenance workflows leading to a reduction in operational costs. IRS server having visibility into both Optical Transport and IP domains will also be able to make optimal use of transport infrastructure.

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 May 7, 2013.

Copyright Notice

Copyright (c) 2012 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

<a href="#">1.</a>	Introduction . . . . .	<a href="#">3</a>
<a href="#">1.2.</a>	Acronyms . . . . .	<a href="#">4</a>
<a href="#">2.</a>	IRS Agent Enhancements . . . . .	<a href="#">4</a>
<a href="#">2.1</a>	Routing changes to degraded optical paths . . . . .	<a href="#">4</a>
<a href="#">2.2</a>	Fast path routing changes to degraded paths . . . . .	<a href="#">5</a>
<a href="#">2.3</a>	Simplified management . . . . .	<a href="#">5</a>
<a href="#">3</a>	Interface to transmission system. . . . .	<a href="#">6</a>
<a href="#">3.1</a>	Direct Mode . . . . .	<a href="#">6</a>
<a href="#">3.2</a>	Indirect Mode . . . . .	<a href="#">7</a>
<a href="#">4.</a>	Conclusion . . . . .	<a href="#">8</a>
<a href="#">6.</a>	Acknowledgements . . . . .	<a href="#">8</a>
<a href="#">7.</a>	IANA Considerations . . . . .	<a href="#">8</a>
<a href="#">8.</a>	Security Considerations . . . . .	<a href="#">8</a>
<a href="#">9.</a>	Normative References . . . . .	<a href="#">8</a>
	Authors Addresses . . . . .	<a href="#">9</a>

## **1. Introduction**

The following is a use case for IRS as a dynamic, multi-layer, policy-based routing service. IP routers have very sophisticated policy and routing capabilities that allow them to manage aggregate traffic flows dynamically based on changing network conditions. But the state of transport paths tends to be viewed by routing systems in a binary fashion as being up or down. Thus, despite the policy and routing intelligence of the router, the decision it makes about a transport path is often very simple: to route traffic over it or to route traffic away from it.

But a transport path may be in a state where it is degraded but still capable of bearing traffic. Consider the case of a path with a high, pre-FEC BER. The FEC is able to eliminate bit errors from reaching the router but the Pre-FEC BER may be predictive of an imminent failure. If a failure does occur, flows for applications using the link will be impacted until a recovery occurs and packets may be delayed or dropped. In such a case, the best policy may be to preemptively route away from the degraded path those flows that emanate from loss-sensitive and delay-sensitive applications. In this way, if the failure does occur, loss and delay sensitive applications will not be impacted. Alternatively, it may be advantageous to raise the IGP or path computation cost metric of the degraded path. In this way, traffic will be routed away from the degraded path unless no other path is available.

In many cases, an application with visibility of both IP and Optical domains through an IRS server could make better decisions about how to manage traffic flows traversing the network.

The goal of this use case is not to specify an all-encompassing system that would replace the NMS or craft interface of the transport system as well as the CLI/Netconf/SNMP interface and IP control plane of the router. The goal is rather to provide a means to automate typical workflows that operate over a router and transport system without interfering with the existing management and control plane systems.

These workflows would be facilitated by a standardized IRS interface. Traditional network/element management systems and methods including TL1, SNMP, Netconf/CLI and proprietary interfaces may still be used for exceptional conditions. In this way, the bulk of operational tasks can be automated over a range of vendor equipment and the maintenance burden may be significantly reduced. Perhaps more importantly, the IRS server can orchestrate complex, policy-based changes to flow paths as a result of changes in the optical domain. Optical paths may be run more efficiently and with lower margin

without propagating route flaps in the IP domain. Similarly, IRS could coordinate dynamic and on-demand coordination of routing changes to provisioning of optical paths.

## **1.2. Acronyms**

BER	Bit Error Rate
CD	Chromatic Dispersion.
CLI	Command Line Interface.
OSNR	Optical Signal to Noise Ratio.
PMD	Polarization Mode Dispersion.
TL1	Transaction Language 1.

## **2. IRS Agent Enhancements**

IRS agents have been proposed as a means to programmatically inject routes into routers and to stream information from them. It would be useful to also have the IRS agents on transport equipment such as transponders, wave selective switches, optical cross connects and amplifiers. The IRS agents would provide data to the IRS server and on to an application. The application could make intelligent decisions about how to direct flows by using realtime information about the state of the optical domain. Basic troubleshooting and provisioning between the router and the transmission equipment could be accomplished through the same programmatic. In this way, IRS would support automation of management, recovery and dynamic bandwidth allocation.

### **2.1 Routing changes to degraded optical paths**

Consider the case of an external transponder and a router interface. The line side may have a variety of parameters showing that the optical path is either degraded or in the process of becoming degraded. The transponder will have indications of decreasing OSNR due to CD, PMD, or even non-linear errors. With post-processing the BER may be kept low enough to avoid the line being taken down. But an increasing error rate may indicate a likelihood of imminent failure. Under this assumption, there are several different ways that the IRS server could react: the IRS server may be able to shift traffic away from such a link before it fails and thus avoid packet loss; It may also be useful for the IRS server to shift critical flows away from a link that may be likely to fail but allow best effort traffic to run over it; Or it may be acceptable to raise the routing metrics of such a link so that other paths will be preferred if they are available.

### **2.2 Fast path routing changes to degraded paths**

Because IRS can install state in both the router and the transmission domain, it is can be used to set up compatible behavior between transmission OAM and the routing system. For example, in the example above, assuming a grey interface between the router and transmission system, IRS could map changes in metrics referring to optical impairments to G.709 or Ethernet OAM indications to be propagated back to the router. The router would then have predefined actions to react most appropriately to the flavor of fault that the transmission network element has been configured to monitor.

These examples are relatively simplistic. It is possible that the correct behavior is much more complex than has just been described. As with many softwaredefined networking initiatives, the real applications will ultimately be developed after the mechanisms are available. The key is that IRS can use information about the optical domain to influence traffic in the packet domain based on the policies of the network administrator.

This model can also be extended to include basic provisioning tasks. It is possible to use IRS to control basic path parameters such as channel assignments and power levels. IRS could also provide basic information about the transponder such as the state of the each of the channels and the presence of alarms. In this way, IRS could provide a basic management interface that could be used to provision and monitor the transponder.

Putting together this management plane integration with the visibility between the routing and optical domain, the router and external transponder would have many of the advantages of a transponder integrated in a router. Some of the important benefits of actual integration- unified management and faster recovery- would be provided by means of this virtual integration method. At the same time, the operator would be able to choose the most appropriate combination of router/switch and optical transmission gear based on cost, technology curve, availability, etc.

In the case where a router is supporting "black links" as per G.698.2, IRS may also provide useful means to extract link property parameters from network element and to correlate them for a path. Although, this is not the primary motivation for extension of IRS to transmission networks, it is also worth considering.

### **2.3 Simplified management**

In transport networks today, a variety of interfaces are supported for the provisioning and management of the network elements including TL1, SNMP, CLI and other standard and proprietary methods. The diversity of these protocols raises the burden of software development for automation tools especially in networks where the equipment from multiple vendors is installed. Yet replacing these

methods is unrealistic.

What is needed is a common interface for a subset of functions related to common workflows. This interface can be standardized across heterogeneous products including routers, switches, transponders, cross-connects, and amplifiers. IRS is a good candidate for this.

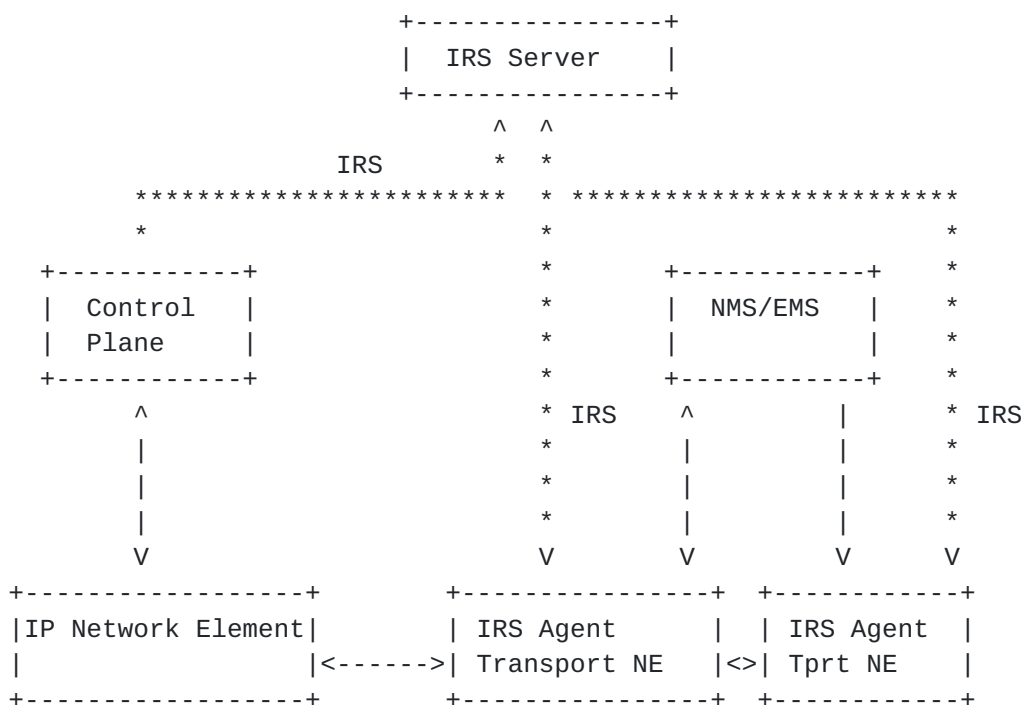
An IRS Agent could manage simple operations on the switch in a consistent way.

### **3 Interface to transmission system.**


There are three modes of access between IRS and the transport system. These three modes are called Direct Mode, Indirect Mode and Hybrid Mode. These modes correspond to whether the IRS agent is logically integrated with the NMS/EMS or the transmission network element or both. In direct mode, the IRS agent is logically associated with the network element. The IRS server can access the network element directly. In indirect mode the IRS server has access to the NMS/EMS system of a transport network. It is possible for either direct or indirect mode to be used independently. Hybrid mode refers to the case where both direct and indirect modes are used simultaneously.

#### **3.1 Direct Mode**

In this mode, the IRS server communicates directly with transport network elements. IRS bypasses whatever NMS or EMS is in place. An IRS agent resides on the transport network element. This method has the advantage of allowing simple and direct access to the Transport Network Elements. Proprietary NMS/EMS systems may still be used and do not need to be integrated into the IRS architecture. Problems need to be considered such as how the NMS/EMS system is synchronizes with changes that committed by the IRS server.



In this mode, IRS server interfaces with an EMS or NMS system. The IRS agent functionality resides within the NMS/EMS and serves as a proxy to the network elements. This mode has the advantage that the NMS/EMS can abstract or virtualize the transport elements. The NMS/EMS system can easily remain synchronized with the IRS server.



#### 4. Conclusion

IRS is a promising means to provide a programmatic interface to routing systems. By extending this capability to transmission networks, IRS would be even more useful. Common maintenance tasks may be automated and network control can be exerted with a view of both IP and Optical domains. As advances in optical networking such as coherent detection and post-processing increase the agility of these networks, IP-integrated, programmatic control and visibility will become even more useful.

#### 6. Acknowledgements

The authors wish to thank Shane Amante for his contributions to this document.

#### 7. IANA Considerations

This document does not raise any particular IANA considerations.

#### 8. Security Considerations

This document does not by itself raise any particular security considerations.

#### 9. Normative References

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

Authors Addresses

Geoffrey Mattson



Juniper Networks  
1194 N. Matilda Ave, Sunnyvale, CA,

Email: [gmattson@juniper.net](mailto:gmattson@juniper.net)

Tom Nadeau  
Juniper Networks  
1194 N. Matilda Ave, Sunnyvale, CA,

Email: [tnadeau@juniper.net](mailto:tnadeau@juniper.net)