

Network Working Group  
Internet Draft  
[draft-ietf-iporpr-framework-01.txt](#)  
Expiration Date: December 2001

Albert Herrera  
Lantern Communications

Russ White  
Cisco Systems

Pankaj K. Jha  
Cypress Semiconductor

Raj Sharma  
Sanjay Agrawal  
Prasad Jogalekar  
Arun Sastry  
Luminous Networks

Khaled Amer  
AmerNet

June 2001

A Framework for IP over Resilient Packet Rings  
<[draft-ietf-iporpr-framework-01.txt](#)>

Status of this Memo

This document is an Internet-Draft and is in full conformance with all provisions of [Section 10 of RFC2026](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at  
<http://www.ietf.org/ietf/1id-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at  
<http://www.ietf.org/shadow.html>.

## Abstract

This document discusses technical issues and requirements of running IP over Resilient Packet Rings. It is the intent of this document to produce a coherent description of all significant approaches, which were and are being considered by the IPORPR working group. Selections of specific approaches, making choices regarding engineering tradeoffs, and detailed protocol specification, are outside of the scope of this framework document.

## Specification of Requirements

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](#).

## Table of Contents

<a href="#">1</a>	Introduction .....	<a href="#">3</a>
<a href="#">1.1</a>	Terminology .....	<a href="#">3</a>
<a href="#">1.2</a>	Overview of IEEE 802.17 RPR .....	<a href="#">3</a>
<a href="#">2</a>	IPORPR Objectives.....	<a href="#">5</a>
<a href="#">3</a>	Motivation for IPORPR .....	<a href="#">5</a>
<a href="#">3.1</a>	<b>Benefits Related to Traditional Broadcast Media Approach.</b>	<a href="#">6</a>
<a href="#">3.2</a>	Benefits Related to a Client/Service Interface to RPR....	<a href="#">6</a>
<a href="#">3.3</a>	Benefits Related to Enhanced L3 Awareness .....	<a href="#">7</a>
<a href="#">4</a>	IPORPR Input and Proposals to 802.17 RPRWG .....	<a href="#">9</a>
<a href="#">4.1</a>	No Changes to IP in a L2-Broadcast-Media Mode .....	<a href="#">9</a>
<a href="#">4.2</a>	Support for a Service Interface Capability .....	<a href="#">9</a>
<a href="#">4.2.1</a>	Proposed Operations .....	<a href="#">10</a>
<a href="#">4.2.2</a>	Granularity of Provisioned BW/Path/Tunnel/Circuit ....	<a href="#">11</a>
<a href="#">4.3</a>	Support for Enhanced Layer3 Operation .....	<a href="#">11</a>
<a href="#">4.3.1</a>	Discovering Logical Adjacencies on the Ring .....	<a href="#">11</a>
<a href="#">4.3.2</a>	Ring Direction Indication .....	<a href="#">12</a>
<a href="#">4.3.3</a>	Fault Interaction .....	<a href="#">12</a>
<a href="#">4.3.4</a>	Traffic Engineering .....	<a href="#">12</a>

<a href="#">4.4</a>	Network Management Support .....	<a href="#">13</a>
<a href="#">5</a>	Acknowledgements .....	<a href="#">13</a>
<a href="#">6</a>	Security .....	<a href="#">13</a>
<a href="#">7</a>	Full Copyright Statement .....	<a href="#">14</a>
<a href="#">8</a>	References .....	<a href="#">14</a>
<a href="#">9</a>	Author's Addresses .....	<a href="#">15</a>
	Appendix A: Logical Adjacency TLV Format and Processes .....	<a href="#">16</a>

Herrera, et al. Expires December 2001 [Page 2]

INTERNET-DRAFT A Framework for IP over RPR June 2001

## [1.](#) Introduction

This document takes into consideration the parallel efforts that are progressing within the IEEE and IETF. IEEE's 802.17 RPRWG is expected to develop an RPR (Resilient Packet Rings) MAC standard that optimizes packet transport over rings on LAN, MAN and WAN topologies. The primary goal of IETF's IPORPR WG is to articulate a set of features that integrates the RPR MAC functions as defined by the [802.17 RPRWG with network layer routing](#).

This document is based on current 802.17 RPRWG objectives and assumptions, which are by no means final.

### [1.1](#) Terminology

- Link Layer or Layer 2 or L2: Refer to the data-link layer such as IEEE 802.3 / Ethernet and, in the case of this document, IEEE 802.17/RPR.
- Layer 2 or L2 devices: Refer to devices that only implement Layer 2 functionality.
- Layer 3 or L3: Layer 3 of the ISO 7 layer model and the use of the Internet Protocol (IP) at this layer.
- Layer 3 or L3 devices: Refer to devices that implement Layer 3 functionality.
- The reader is assumed to be familiar with the terminology as defined in [\[1\]](#) [\[2\]](#) [\[3\]](#) [\[4\]](#) [\[5\]](#).

### [1.2](#) Overview of IEEE 802.17/RPR

Traditional metro and core networks were designed and optimized for voice and video using the SONET/SDH circuit switch layer. Today's dominant traffic sources are data applications that drive bandwidth demands in both the MAN and WAN. Carrying data traffic on traditional circuit switched networks has proven to be inefficient, complex and expensive.

The IEEE 802.17 Working Group was established to formulate a new Resilient Packet Rings (RPR) MAC layer standard. RPR is designed to be a packet switched technology optimized for data traffic while maintaining carrier class operational attributes previously available only through circuit switch services.

Herrera, et al.	Expires December 2001	[Page 3]
INTERNET-DRAFT	A Framework for IP over RPR	June 2001

Some objectives adopted by the 802.17 RPRWG include:

- Support for a dual counter-rotating ring topology
- Support for destination stripping of unicast traffic
- Support for protection switching in less than 50ms
- Support for data rates of (at least) 1 Gb/s to 10Gb/s
- Support for a MAC that is both PHY and payload agnostic
- Support for a fully distributed access method (no master node)
- Support for multi-vendor interoperability at the ring level
- Support for multicast traffic
- Support for managed objects or management "hooks" required for efficient operational control of the network.

These objectives define a network topology made up of closed-loop, point to point, spans with the MAC layer creating the resulting logical ring topology. At the physical layer, RPR looks like a set of point-to-point spans while at the datalink layer it appears to be a broadcast media LAN (equivalent to Ethernet, for instance). The dual logical rings created are counter-rotating, carrying both control and data packets. Packets are destination stripped for efficient bandwidth utilization and prioritized for appropriate treatment during protection situations. Protection mechanisms will meet or better the SONET's 50ms service guarantees without the need for dedicated or reserved bandwidth for redundancy. Data rates are defined at 1 Gb/s to 10Gb/s but do not preclude support for lower or higher rates in the future.

These are initial objectives passed by the 802.17 WG with more to follow. IPORPR, further described in [section 2.0](#), will build upon this set of features.

This initial set of objectives provides RPR with the following key attributes.

#### Bandwidth Efficiency:

Unlike traditional SONET networks that require as much as 50% of ring bandwidth for redundancy, RPR utilizes both dual counter rotating rings for control and data traffic while still maintaining SONET APS-like protection mechanisms. Spatial reuse is achieved by allowing a bi-directional flow of traffic on ring spans between source and destination nodes. The destination node strips unicast packets from the ring. When a packet is "stripped" from the ring, it no longer consumes bandwidth on the ring, but frees up downstream spans to be utilized by other packets.

Herrera, et al.	Expires December 2001	[Page 4]
INTERNET-DRAFT	A Framework for IP over RPR	June 2001

#### Protection:

RPR's objective is to provide 50ms service protection similar to SONET APS. Approaches under consideration are to wrap and/or steer traffic away from the fault. When 'wrapping', nodes adjacent to the fault will 'wrap' traffic around onto the other ring (i.e. inner ring traffic to outer) allowing flows to maintain connectivity to the destination node by way of a long path. 'Steering' will actually divert flows to the destination via the long path. A combination of both is also an option where nodes beside the fault will first 'wrap' then steer at their convenience.

This protection feature will be available while fully utilizing both counter-rotating physical fibers. Traffic priorities ensure appropriate handling of high priority vs. best effort traffic during fault situations.

#### Simplified Service Provisioning:

One RPR objectives is for a distributed access method. This feature together with the fast protection and automatic re-establishment of

services, provide for plug-and-play mechanisms for quick insertion and removal of nodes.

RPR is also designed as a packet switched technology utilizing shared bandwidth within the rings, with nodes aware of available ring capacities. The simplicity becomes evident on full-mesh connectivity requirements where traditional circuit-based models require  $O(n^2)$  point-to-point connections versus RPR's single service connection to the ring.

## **2. IPORPR Objectives**

The primary goal of the IPORPR Working Group is to articulate a set of features and functionality that will allow IP to integrate smoothly with the RPR MAC functions as described in [Section 1.2](#). This will be presented as input to IEEE's 802.17 RPRWG for consideration. The intent is to align the needs of the Internet and IP with the emerging 802.17 standard with the hope that the 802.17 RPRWG will make design choices that take these requirements into consideration and in some instances, avoid duplication of efforts.

## **3. Motivation for IPORPR**

Network layer routing can treat these RPR rings as a broadcast capable LAN media similar to Ethernet without requiring any change to current L3 protocols. With minimal change, it could also leverage RPR's unique features, such as dual-counter-rotating rings, destination stripping (which effectively provide two possible paths from a source node to a destination node) or 50ms protection switching (capable of masking any single physical fault).

Herrera, et al.	Expires December 2001	[Page 5]
INTERNET-DRAFT	A Framework for IP over RPR	June 2001

There are three basic approaches to L2/L3 interaction presented in this document (further described in the following sections):

- L3 view of an RPR ring as a traditional broadcast media
- A L3 service interface to an RPR ring for specific resource/services requests
- Enhance L3 awareness and interaction with an RPR ring.

Introducing additional L2/L3 interaction will provide increased ability to better utilize ring resources and increase efficiency and scalability at Layer 3.

The following are projected benefits related to IPORPR.

### **3.1 Benefits Related to Traditional Broadcast Media Approach**

In an approach where L3 treats an RPR Ring as a traditional data link, broadcast media (e.g. Ethernet), L3 will rely on L2 to provide the benefits of the underlying features such as protection mechanisms, congestion control and bandwidth management within each RPR ring domain.

This is the simplest approach to L2/L3 interaction and requires no change whatsoever to existing L3 protocols.

### **3.2 Benefits Related to a Client/Service Interface to RPR**

RPR rings include a distributed access management capability that allows rapid establishment and reconfiguration of flows and/or connections, reducing provisioning times and effectively lowering costs while providing the means to set and guarantee SLAs and QoS configured on either a per-flow and/or connection basis. The projected new services for RPR rings are bandwidth on demand, point and click establishments of circuits/paths and VLANs and/or VPNs.

A standardized interface between RPR and the service layers such as IP will allow end-to-end internetworking of paths and flows through these rings and will provide the benefits of RPR end-to-end even if other networks are involved (i.e. meshes).

This service interface will allow clients to the RPR ring to signal parameters (i.e. SLA constraints, bandwidth requirements, etc.) to L2 for specific path selection or other L2 behavioral characteristics. The signaling and routing interface between the client (IP and others) and the RPR ring is similar to traditional User-Network Interface (UNI).

This allows either a dynamic and/or static set of capabilities for service provisioning without closely coupling the underlying RPR ring operations to the upper layers.

Herrera, et al. Expires December 2001 [Page 6]

INTERNET-DRAFT A Framework for IP over RPR June 2001

This interface can support a range of clients from individual customers to lower tier ISPs connected to a carrier's carrier for dynamic provisioning of services or the service provider's management console for static establishment of services.

### 3.3 Benefits Related to Enhanced L3 Awareness

Enhanced L3 awareness of the underlying physical topology will allow greater L3 controls of certain ring operations/behaviour such as choice of ring direction, alternate fault response, traffic engineering and related QoS metrics.


Note that this document does not address issues related to enhanced Layer 3 awareness across bridged, multi-ring domains. In topologies such as these, it might be best for Layer 3 to simply treat these as simple broadcast domains.

The following are related L3 components that can be enhanced when running IP over RPR rings.

#### Cost Metrics Awareness:

At the MAC layer, RPR may implement its own topology discovery mechanism within a single set of dual-ring implementation. This may involve tracking source and destination MAC addresses and utilizing simple metrics, such as number of hops, to determine the shortest path and most efficient ring direction for delivery of unicast packets.

For instance, in the network depicted below, the cost from node A to node C would be 1 hop across the directly connected side of the ring, while it would be two hops through node B.



However, using L3 associated costs for links between the nodes (based on physical distance or some other link trait, for instance), the path through node B is shorter.

Communicating the preferred ring direction to L2, based on layer 3 metrics can optimize packet delivery - especially in topologies with dramatic differences in node distances and costs.




## Adjacency and Fault Awareness:

Given RPR's destination stripping ability, and once fault notification is enabled between Layer 2 and 3, it would be beneficial for Layer 3 to be aware of directly adjacent neighbors for route recalculation during fault situations.

Within a single ring domain, the benefits of adjacency awareness might not be obvious since the underlying RPR mechanisms will self-heal and send flows efficiently around or away from faults to maintain connectivity. However, L3 route recalculation might be necessary to fully reflect the changed topology during these events especially when the user deploys a mix of ring and mesh topologies. This will allow a recalculation of paths based on established cost metrics.

For instance, in the network depicted below, a ring made up of nodes A, B, C and D is bisected by a P2P link between D and B with a cost of (5). Under normal ring operations, C will reach B in a counter-clockwise direction via a single span with a cost of (5). Should a fault occur on this span between C and B, C will reach B via D skipping A altogether due to an expensive span cost when traversing A on a clockwise ring direction.



Notifying L3 of fault situations from the lower layers will allow swifter response from L3 should this be required by the user. Fault situations that L3 should be aware of are those that would cause ring traffic to be wrapped or redirected away from the fault. Other minor fault situations, where ring operations continue to function, might not have to be communicated to L3.

## TE/Explicit Routing/QoS Routing Support:

Although RPR might incorporate its own topology discovery, bandwidth management and congestion/flow control mechanisms at L2, there might be instances wherein L3 would prefer to override these.

MPLS, used for traffic engineering, offers several benefits within an RPR-ring topology. One benefit is the capability to establish explicit switched paths that are not constrained by any of RPR's forwarding paradigms. The MPLS preferred path along the ring could be the long-path towards a destination node. Another benefit is the capability to associate attributes to these traffic flows and to administer these flows across diverse topologies (combination of

meshes and rings). This same mechanisms will enable L3 per flow QoS treatment in which the path chosen for a particular stream is chosen in response to the QoS constraints.

Herrera, et al.	Expires December 2001	[Page 8]
INTERNET-DRAFT	A Framework for IP over RPR	June 2001

#### **4. IPORPR Input and Proposals to 802.17 RPRWG**

The IPORPR Working Group was chartered by the IESG to produce a requirements and framework document, which can be used as input to the IEEE 802.17 RPRWG. The intent is to help the RPRWG formulate its requirements. Areas of particular interest to the IPORPR WG are: Layer 2 to Layer 3 interaction in: alarm notification; fast restoration; fast convergence; traffic engineering and quality of service issues.

The primary requirement set forward by IPORPR is that the IEEE 802.17 WG makes design decisions that will have the least impact and require the least amount of change to IP.

The following are the proposed requirements for L2/L3 interaction.

##### **4.1 No Changes to IP in a L2-Broadcast-Media Mode**

The first and foremost requirement for 802.17 is to allow IP to operate without any changes, whatsoever, when treating RPR rings as a traditional L2 broadcast media. This approach should not be any different from how IP operates today over other data link layers such as Ethernet. This will allow existing protocols and design constructs to apply unchanged.

##### **4.2 Support for a Service Interface Capability**

A 'service interface' can be thought of as a User-Network Interface (UNI). A connection, in a UNI sense, is defined by its demarcation from the ingress node to egress node. These can be within the same ring or might span other rings or meshes. The behaviors of these pairings are defined by their connection attributes. A standardized set of parameters is required for the edge-facing interfaces. This will include service requirements and attributes during normal ring operations and during fault situations. Specifics for these parameters are not discussed in this current framework but should be developed as the 802.17 standards evolve and as more detailed features and behaviors are firmed up.

There are two approaches in terms of request handoffs from a service

interface.

One approach is to handoff requests to L2 and use L2 capabilities, such as topology discovery, path selection and bandwidth accounting, to service requests from the client interface. RPR's ring operations should be capable of satisfying these requests in terms of contracts associated with SLAs.

Herrera, et al. Expires December 2001 [Page 9]

INTERNET-DRAFT A Framework for IP over RPR June 2001

The other approach is to allow L3 intelligence to meet these requests. In cases where enhanced L3 awareness is enabled for traffic engineering, path establishment, bandwidth accounting, QoS and associated signaling, service parameters can be handed over to L3. This approach will accommodate requirements for extending any existing L3 provisioning and signaling models (i.e. on mesh networks) to and through RPR ring topologies.

Both approaches should be supported.

#### **4.2.1 Proposed Operations**

A client upon ingress to the RPR ring may request:

- a new connection subject to policy
- a tear-down of an existing connection
- attribute modifications to an existing connection based on policy
- attribute queries on an existing connection
- status queries of issued requests

Admission control in terms of accepting or rejecting requests depends on the underlying provisioning and signaling capabilities. Local node policy will dictate when connections may be established and whether these are dynamically established or administratively controlled. Parameters and associated mechanisms are required to validate interoperable capabilities at the ingress and egress ports of a requested connection. Neither the connection setup mechanism nor the method for identifying flows is addressed in this document.

Bounded delay guarantees require that every RPR node in the connection path supports a guaranteed service feature whether within the L2 operations itself and/or via L3 enhancements, resource accounting and policy enforcements. Each RPR node accepting a request for service must ensure that adequate bandwidth and packet processing resources are available to handle the request as specified in the

client's request parameters / Traffic Specifications. This may be implemented through active admission control based on resource availability such as link bandwidth, port buffer space, forwarding engine capacity or transit traffic treatments. The assumption is that [802.17](#) will provide the appropriate architecture to meet stringent delay and jitter guarantees whether through appropriate queuing mechanisms, media access management or transit traffic mechanisms.

In cases where the RPR node detects a failure on the ingress physical port, or if the ingress port is administratively disabled, the corresponding service connection, related mappings and reservations may be withdrawn. Payload synchronization and maintenance alarm requirements are outside the scope of this document.

#### **4.2.2 Granularity of Provisioned Bandwidth/Path/Tunnel/Circuit**

As stated on the objectives, RPR rings will support at a minimum, data rates, on the ring, of 1Gb/s to 10Gb/s with even higher rates for consideration within 802.17 in the future. At the ingress service interface, an upper limit equivalent to the fastest link on the ring should be allowed, as well as lower granularities such as 1Gb/s, 10/100mb/s, 128kb/s, 56kb/s, etc.

#### **4.3 Support for Enhanced Layer3 Operation**

Upper layer protocols will require various interaction points with RPR rings to take effective advantage of the available bandwidth (in both directions) through the ring. Generally, these interaction points will fall into two categories: Informing the data link layer of ring direction choices made by Layer 3 protocols for each next hop along the ring, and informing the Layer 3 protocols of the ring topology.

The following are technical approaches for consideration:

##### **4.3.1 Discovering Logical Adjacencies on the Ring**

IS-IS and OSPF require knowledge of all neighbors on a ring with the additional information of which neighbor is directly connected upstream and directly connected downstream. Since all devices attached to a ring may not be running layer three protocols, simply using the physical topology discovered by the data link layer isn't an effective way to discover all layer three topology information. Instead, it will be more effective for RPR to provide an opaque transport through which routing protocols can determine which neighbors are logically attached upstream and downstream.

To this end, the RPR data link layer should be able to carry 'opaque data' TLV's (Type/Length/Value) between nodes, accepting data from upper layer protocols, and passing this information back up the stack on adjacent nodes only (Refer to [Appendix A](#): Logical Adjacency TLV Format and Processes).

Through this process, devices running layer 3 protocols will discover its directly adjacent downstream and directly adjacent upstream neighbors.

NOTE: This is one implementation possibility used to explain the behavior; this document is not intended to restrict or dictate the correct implementation to achieve this behavior. This could, for instance, be implemented using multicast between the nodes, which

layer two could then compare to the physical topology discovered through other means, passing only the information about 'logical neighbors' up to layer 3, or some other mechanism.

#### [4.3.2](#) Ring Direction Indication

A mechanism **MUST** exist for Layer 3 to override any indication by Layer 2 as to the direction to be followed on the ring to reach any of its nodes. The two possible directions are "upstream" and "downstream". Administrative policies and/or traffic engineering parameters will influence layer 3 decisions.

If Layer 3 does not indicate one of these directions, then Layer 2 **SHOULD** use its own methods to determine it.

Should Layer 2 determine the ring direction, there might be instances where the direction taken by RPR be explicitly indicated to Layer 3. These are instances where strict latency and jitter requirements are to be met and ring direction decisions will require indications from Layer 2.

#### [4.3.3](#) Fault Interaction

If a device on the ring or a link between two nodes on the ring fails, the ring will wrap or steer traffic away from the fault. During these fault situations, nodes adjacent to the fault might want to perform route re-computation and advertise a single ring adjacency in their LSP. A signal **MUST** be provided to the layer 3 protocols on the attached nodes indicating the loss of a logical adjacency. Layer **2 MUST NOT provide information about the newly reachable neighbor** through the wrapped or re-directed ring, as this could cause layer 3 to be unable to properly converge on the new topology.

In cases where a node failure occurs, where a particular device leaves the ring but maintains transit paths and ring connectivity in place, normal RPR ring operations are expected to proceed as normal. Any fault interaction is expected to be resolved at a higher layer where procedures are in place to react a lost destination nodes or adjacencies (e.g. hellos, keepalives, topology refresh, etc.).

#### [4.3.4](#) Traffic Engineering

To take full advantage of RPR, TE capabilities should be made available to manage bandwidth on both inner and outer rings.

Sometimes it is preferred to force a packet to follow an explicitly chosen route different from the path chosen by a dynamic routing protocol or different from the shortest path chosen by RPR. This implies the capability to determine the ring direction a particular flow or tunnel will take (long path or short path) and the capability to establish an explicitly routed path from ingress to egress.





RPR should be able to meet the SLA constraints of established flows that originate from outside the ring whether from another ring or from a mesh topology. This implies the capability to apply resource reservations along a chosen path and the capability to recognize precedence and class of service parameters for applicable discard or scheduling mechanisms.

In terms of MPLS support, different approaches for label stack encoding, to enable forwarding of labeled packets encapsulated within RPR headers, can be referenced in [5].

This document does not suggest use of a single label distribution protocol. Please refer to [4] for options and implications.

This document does not cover interaction between RPR protection switching and MPLS based fault recovery mechanisms.

#### [4.4](#) Network Management Support

RPR MUST support operations, administration, and maintenance facilities that are at least as extensive as those supported in current IP networks. Current network management and diagnostic tools SHOULD continue to work in order to provide some backward compatibility.

### [5.](#) Acknowledgments

The ideas and text in this document have been collected from a number of sources and comments received. We would like to thank John Coulter, Kshitij Kumar, Alvaro Retana, Don Slice, Liem Nguyen, Thierry Paient, Jason Fan, Vinay Bannai for their inputs and ideas.

### [6.](#) Security

Security in a network using IPORPR should be relatively similar to security in a normal IP network. The route filtering and packet filtering features are unchanged.



## **7. Full Copyright Statement**

"Copyright (C) The Internet Society. All Rights Reserved. This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into."

## **8. References**

- [1] ISO 10589, "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)" [Also republished as [RFC 1142](#)]
- [2] [RFC 1195](#), "Use of OSI IS-IS for routing in TCP/IP and dual environments", R. W. Callon, Dec. 1990
- [3] [RFC 2178](#), "OSPF Version 2", J. Moy. July 1997
- [4] [RFC 3031](#), "Multiprotocol Label Switching Architecture", E. Rosen, A. Viswanathan, R. Callon, January 2001.
- [5] [RFC 3032](#), "MPLS Label Stack Encoding", Rosen, E., Rekhter, Y., Tappan, D., Fedorkow, G., Farinacci, D. and A. Conta, January 2001.
- [6] IEEE 802.17 RPR PAR and 5 Criteria, <http://www.ieee802.org/17>, November 2000.



**9. Author's Addresses**

Albert Herrera  
Lantern Communications  
2000-1642 Merivale Road,  
Ottawa, ON K2G 4A1  
Phone: 613-727-7112  
Email: aherrera@lanterncom.com

Russ White  
Cisco Systems, Inc.  
7025 Kit Creek Road  
Research Triangle Park, NC  
Phone: 919 392-3139  
Email: ruwhite@cisco.com

Pankaj K Jha  
Cypress Semiconductor  
3901 N First Street  
San Jose, CA 95134  
Phone: 408 432 7091  
pkj@cypress.com

Raj Sharma (raj)  
Sanjay Agrawal (sanjay)  
Prasad Jogalekar (prasad)  
Arun Sastry  
(arun@luminousnetworks.com)  
Luminous Networks  
10460 Bubb Road  
Cupertino, CA 95014  
Tel: 408-342-6400

Khaled Amer  
AmerNet  
Email: khaledamer@usa.net



## Appendix A: Logical Adjacency TLV Format and Processes

IS-IS and OSPF require knowledge of all neighbors on a ring with the additional information of which neighbor is directly connected upstream and directly connected downstream. To this end, the RPR data link layer should be able to carry 'opaque data' TLV's (Type/Length/Value) between nodes, accepting data from upper layer protocols, and passing this information back up the stack on adjacent nodes only.

So, for instance, in this small sample ring, one possible implementation of this functionality would be:



Assume that A, B, and C are running some layer three protocol, while D is not. A needs to know which devices on the ring are possible next hops upstream and downstream. To facilitate this:

- A, B, and C will inject information into a Logical Adjacency TLV (LAT) which will be carried by the RPR data link layer both upstream and downstream.
- When A receives B's LAT, the RPR data link processes will determine which protocol stack to pass the opaque information within the TLV to, and pass it up the stack with an indication of which direction is used to reach this adjacent device.
- When D receives C's LAT, it will discover that it has no layer three protocol running, so it will simply forward the LAT on to A.
- When A receives C's LAT, the RPR data link processes will again determine which protocol to forward the opaque portion of the data to, and pass it up the stack with an indication of which ring direction is used to reach this adjacent device.

Through this process, A will discover that it's directly adjacent downstream (clockwise) neighbor is B, while it's directly adjacent upstream (counter clockwise) neighbor is C. A's routing processes can use this information, in conjunction with information learned at layer three, to make optimal routing decisions through the ring.





### Logical Adjacency TLV Format

```
-----  
| src mac| id_number | id_type | id_len | id | ...  
-----
```

- o src mac: The MAC address of the source ring node
- o id\_number (1 byte): The number of this identifier within the packet
- o id\_type (1 byte): described below
- o id\_len (2 bytes): length of the opaque information contained in the id section
- o id: The actual opaque information provided by the layer 3 protocol, and passed to the receiving device's layer 3 protocol.

### ID Type Field

While the information contained within the logical adjacency TLV's is opaque to the RPR data link layer, the TLV's themselves need to be identifiable in some way so the layer 3 protocol on one device receives the correct TLV from the layer 3 protocols on connected devices. To this end, the TLV should be encoded with a protocol identifier of some type.

Three protocol identifiers are defined here, and others can be identified as needed in the future.

- o OSPF: 1
- o IS-IS: 2
- o BGP: 3
- o EIGRP: 4

