# RPR—Building A Better Ethernet?

Robert Bellman

## Will metro carriers go for resilient packet ring technologies? Will vendors be able to agree on standards?

Network service providers need a new kind of optical ring technology to deliver high-speed metro Ethernet services—at least that's what some vendors think. Companies like Appian Communications, Cisco, Dynarc and Lantern Communications are developing resilient packet ring (RPR) technologies, which replace Ethernet's media access control (MAC) layer with a new ring-based MAC. Atrica and a few others are addressing the challenge with ring-based solutions that preserve the Ethernet MAC. All these approaches promise service providers the best of both worlds: Ethernet's low cost and packet data efficiency, plus SONET's ring structure, reliability and rapid restoral.

Sensing a groundswell, the IEEE formed the 802.17 working group late last year to standardize RPR. A marketing consortium, the RPR Alliance, was formed in January. The IETF also has a working group called IP over RPR (IPoRPR).

But 802.17 has a sweeping agenda that could slow progress toward its most basic goals. A recent RPR Alliance press release declares, "RPR will support carrier-class, service-level-agreement-based metro Ethernet, IP, and legacy TDM services." If I were a member of the 802.17 working group, the breadth of that statement would worry me. The working group needs to keep its eye on the ball—efficient, robust metro-area data services—or 802.17 is destined for a one-way trip to the protocol museum.

### What Goes Around May Not Come Around

Service providers like their existing SONET rings—and for good reasons. Besides SONET's 50-millisecond restoral, the rings reach more customers with less fiber than other topologies and require fewer switch ports at busy hub sites. But SONET, which was designed for voice circuits, wastes bandwidth when carrying packet traffic, and the spanning tree algorithm that guides most Ethernet switches purposely breaks rings in order to prevent bridging loops.

Packet over SONET (POS) offers a partial solution, but only for point-to-point links. A better solution, according to the RPR proponents, would be a new packet MAC that uses rings efficiently but also exhibits the resilience and QOS of SONET.

The 802.17 working group plans to have a first draft standard in January 2002, with final technical changes by September 2002 and official standardization by March 2003. A healthy mix of service providers and equipment manufacturers attends the working group meetings. As usual, each equipment manufacturer is lobbying for a standard that fits its technical vision (and minimizes its re-engineering). The working group has

**FIGURE 1 Destination stripping**

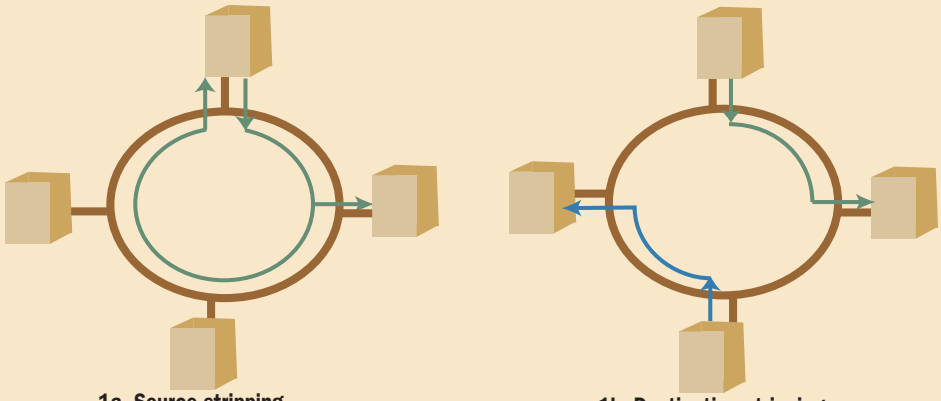

1a Source stripping

1b Destination stripping

*Bob Bellman is president of Brook Trail Research (www.brooktrail.com), a consultancy specializing in network technology assessment and marketing.*



Use BCR's Acronym Directory at www.bcr.com/bcrmag

used these inputs to build consensus for high-level objectives before diving into the technical details.

These objectives, as reported by the RPR Alliance, comprise a rough outline of the architecture, commitments to motherhood and apple pie, and some noticeable opportunities to stray from the basic mission:

- **Dual counter-rotating ring topology:** In dual-ring topologies, SONET uses only one ring to carry live traffic; the other is reserved as a backup. To increase fiber utilization, RPR will send traffic over both rings (in opposite directions) during normal operation.
- **A fully distributed access method without a master node:** An RPR ring will continue operating despite the loss of any node.
- **Protection switching in less than 50 milliseconds:** In the event of a fiber break or node failure, RPR will restore service at least as fast as SONET. This will likely entail linking the two rings into one, so some traffic may get bumped.
- **Destination stripping of unicast traffic:** In some older packet ring architectures, the source node removes unicast packets after they come all the way around (Figure 1a). With RPR, destination nodes remove their unicast packets, freeing downstream bandwidth for reuse by other flows (Figure 1b). Together with packet multiplexing and counter-rotating rings, destination stripping more than doubles RPR's total throughput compared to SONET.
- **Support for multicast traffic:** Multicast packets travel once around the ring to reach every node. In contrast, mesh networks must replicate multicast packets in order to reach all destinations.
- **Support for up to 10 Gbps**: RPR will be fast enough to carry Gigabit and 10-Gigabit Ethernet, but will also support lower data rates.
- **Support for both SONET/SDH PHY (physical layer) and the 1- and 10-Gbps Ethernet PHYs:** Support for existing PHYs will let RPR products use widely available components.
- **PHY and payload agnostic:** To be truly universal, the RPR MAC will be independent of the PHY layer and will not interfere with customer payloads.
- **Plug-and-play support:** New nodes may join the ring without manual configuration.
- **Managed objects:** By defining managed objects, the RPR standard will facilitate OSS integration.
- **Support for services that require bounded delay and jitter and guaranteed bandwidth:** RPR will deliver TDM-like QOS. Welcome to the top of the slippery slope. If RPR strays too far from Ethernet in attempting to support TDM, it won't be able to use commodity chips, and it may pose more direct competition to SONET than some suppliers would like.
- **Dynamic weighted bandwidth distribution:** RPR will allocate bandwidth to competing flows on demand.
- **Support for multiple service types:** RPR will adapt to future requirements.
- **Vendor interoperability:** RPR equipment from different vendors will interoperate on the same ring.

In other words, RPR's primary mission is to make optical rings more efficient for packet traffic, but the standardization effort has attracted some developers who want to see it do more, especially in the realm of QOS and traffic control. Others ask why Layer 3 mechanisms (such as DiffServ and MPLS) can't be used to provide these functions.

**Some metro-area Ethernet vendors are opposed to the new RPR MAC**

### If I Had a Hammer

Another opportunity to go astray is represented in the 802.17 working group's mission statement: "The IEEE 802.17 Resilient Packet Ring Working Group will define a Resilient Packet Ring Access Protocol for use in Local, Metropolitan and Wide Area Networks for transfer of data packets at rates scalable to many gigabits per second."

The last protocol that tried to cover "Local, Metropolitan and Wide Area Networks" was ATM. Untold millions of dollars were wasted trying to make ATM, which was initially conceived as a wide-area technology, suitable for local-area networks.

Is 802.17 making a similar mistake? Bob Love, Chair of the RPR Alliance and Vice-Chair of the 802.17 working group, thinks not. "If you have a hammer, the world looks like a nail," said Love, "but it's clear that MANs are the sweet spot." For the sake of RPR's original mission, let's hope the rest of the working group agrees.

More substantive concerns center around the protocol "sandwich" RPR will create. Users will see an Ethernet service interface that plugs directly into their routers and switches. Underneath that interface, the RPR MAC will move packets around the ring, while underneath the MAC, an Ethernet PHY formats the bits onto the fiber.

Some metro-area Ethernet developers are opposed to the new MAC altogether, including David Yates, Atrica's marketing VP. "RPR is as similar to Ethernet as token ring is to Ethernet; that is, they're not similar at all," he said. "Ethernet is standard, understood and based on cheap components that scale fast. RPR will need different components, different management, and so on." He asserted that RPR-based Ethernet will cost more than ordinary Ethernet.

Atrica takes a different approach with its family of Optical Ethernet Switches. It uses a 10-Gbps Ethernet MAC over WDM wavelengths and relies on existing standards to make up for Ethernet's shortcomings. For example, DiffServ and 802.1p let switches manage traffic priorities, while Multiprotocol Label Switching (MPLS) provides fast recovery from outages.

Meanwhile, the folks at Appian aren't entirely opposed to an RPR MAC, but they also are con-

**Cisco's Spatial Reuse Protocol meets many RPR objectives**

cerned about its complexity. Anand Parikh, Appian's vice president of product marketing and business development, agrees that matters like fairness and QOS should be handled by existing standards outside of RPR. "Keep the architecture simple," he said. "Otherwise, nothing will get done. RPR should use existing standards where possible. There's no need to reinvent the wheel."

Appian is also concerned that a complex MAC layer will pollute the PHY layer and make low-cost Ethernet framers unusable. If RPR can't deliver Ethernet services at Ethernet prices (or better), who will pay for the equipment?

### Yield To Oncoming Traffic?

By contrast, Cisco doesn't share Appian's or Atrica's concerns. In fact, it already has a new MAC layer, the Spatial Reuse Protocol (SRP), that meets many of the RPR objectives. SRP is used in Cisco's Dynamic Packet Transport (DPT) product line, which has been shipping for more than two years, largely to cable modem service operators. Cisco submitted SRP to the IETF in 1999 as informational RFC 2892, and participates in both the IEEE working group and the RPR Alliance.

Cisco's first DPT products were router interfaces, and the initial applications were for ISP intra-POP rings. "We got a head start," said Jeff Baher, a senior marketing manager at Cisco, "before metro Ethernet services were such a big deal." Compared to an ATM or POS mesh, DPT gave service providers a way to keep OC-3 port counts low, and to minimize the number of IP subnets within a site. From there, DPT spread to inter-POP connections and use by cable modem operators to connect head-ends with data centers. Now Cisco boasts more than 12,000 installed DPT ports at more than 160 customers, and Baher expects DPT "to migrate to our high-end Ethernet switches, especially our IP-oriented ones."

When planning that migration, however, Baher is concerned about where 802.17 compliance might lead. In particular, Baher questions the relationship between packets entering the RPR ring and those already on the ring (Figure 2). Should transit packets take precedence over entering packets, or should all packets compete equally for bandwidth at every hop? So far, the 802.17 working group has not decided.

With Cisco's SRP, transit packets take precedence over entering packets, so there's no packet loss on the ring itself. "This is different from how Ethernet switches work," Baher pointed out. Ordinary Ethernet switches, which lack a cut-through path for transit traffic, exhibit varying packet loss throughout the network as traffic congests at each node. If the 802.17 working group decides against transit cut-through, then RPR will lose a potential advantage over ordinary Ethernet switching.

Transit traffic should take priority, agrees Bob Schiff, senior director of strategic marketing at Lantern Communications, and that should be part of the standard. "You must protect the traffic that's already in the ring," he said. "It's like a traffic rotary. Once you're in the rotary, you're okay. The only challenge should be insertion delay." (I guess Schiff hasn't driven in Boston lately. We accelerate into rotaries.) Lantern's Metro Packet Switch products, which employ a cut-through scheme, are scheduled for field trials this fall.

### SONET Or Simple?

Another controversial topic is the relationship between RPR and TDM. At stake are not only the definition and scope of RPR, but also the migration path from SONET. "There are two camps," said Cisco's Baher. "One sees RPR as a SONET ADM replacement. The other sees RPR as simply a high-performance packet MAC."

The rationale of the SONET replacement camp echoes every convergence vendor's slideware: Voice is a big revenue producer, so RPR has to handle voice efficiently. In contrast, the keep-it-simple camp argues that TDM support will be RPR's undoing.

"There's no need to build the MAC of all MACs," insisted Baher, "no need to pull out all the SONET gear. RPR is okay for IP phones, but not for TDM *per se*. If RPR tries to do everything, it won't do enough TDM for TDM folks, and it won't do enough packet for packet folks. Don't kill RPR by overloading it."

Appian's marketing VP, Karen Barton, took the simplicity argument one step further. "It's about today's reality," she said. "RPR needs to coexist with traditional SONET mechanisms. When TDM is fully packetized, the need for SONET will go away. But how many big service providers have accepted packets for carrying voice or private line?"

Appian, therefore, offers a hybrid approach. Its Optical Service Activation Platforms (OSAPs) dedicate some channels on a SONET ring to packet traffic and leave others for ordinary TDM. This allows OSAPs to share fiber rings with conventional SONET ADMs.

**FIGURE 2 Competition For Ring Access**

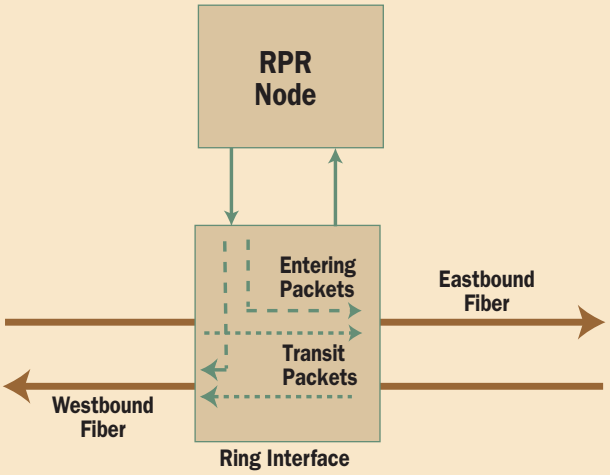

**FIGURE 3 Atrica Optical Ethernet Switches**

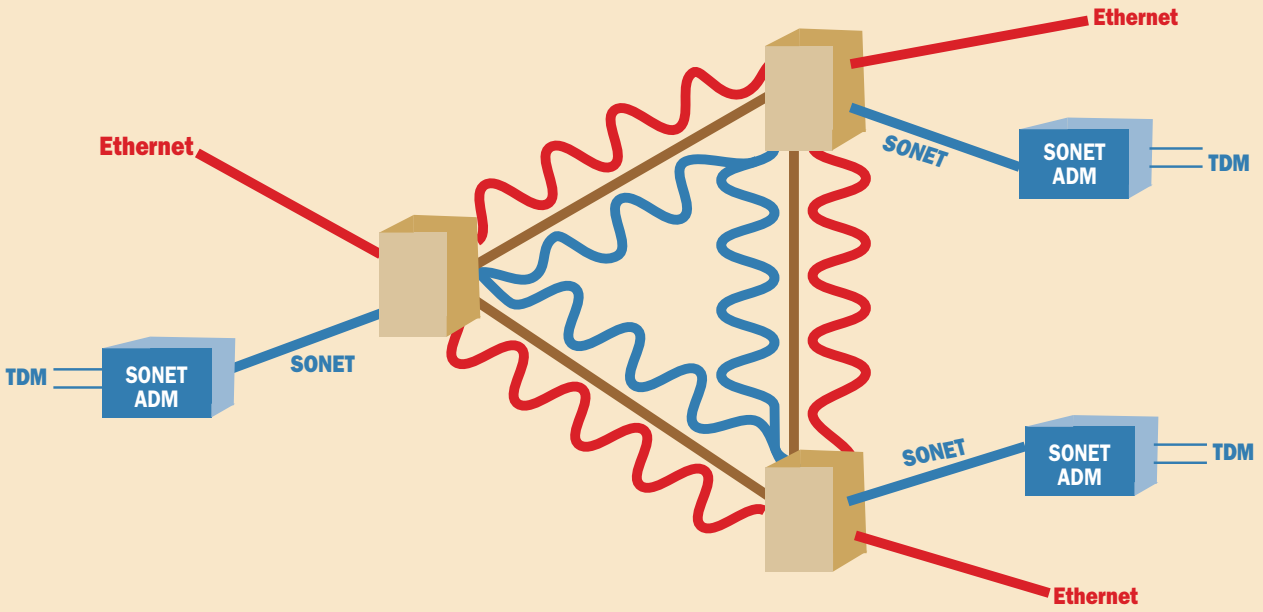


Dynarc's implementation is more a blend than a hybrid. Dynarc is headquartered in Sweden, where "networks are a couple of years ahead of the U.S.," according to Fredrik Hanell, the company's vice president of marketing. Dynarc is currently supporting a project in Stockholm that is bringing 100-Mbps Ethernet-based network access to 20,000 apartments.

Dynarc's Channelized Reserved Services (CRS) architecture integrates IP with Dynamic synchronous Transfer Mode (DTM), an ETSI standard for on-demand circuit switching. CRS routers carry everything in IP packets and group the packets into channels. "A channel is a logical resource defined by anything in an IP packet," explained Hanell, "source address, destination address, QOS bits, whatever."

Channels, in turn, are mapped into TDM-like timeslots. Channel capacity can be fixed, perhaps to tunnel a T1 through the network, or flexible to accommodate less sensitive flows. One DTM feature that Hanell would like to see in RPR is the ability to expand ring capacity by turning on WDM wavelengths.

Yet another variation on packet/circuit coexistence is offered by Atrica's family of Optical Ethernet Switches. As mentioned above, Atrica rejects the RPR MAC in favor of a 10-Gbps Ethernet MAC over WDM wavelengths. Moreover, Atrica's Optical Ethernet Switches carry low-speed TDM traffic in Ethernet frames. At speeds from T1 (1.5 Mbps) to OC-12 (622 Mbps), Atrica uses Ethernet circuit emulation. At those speeds, serialization delay is not a problem, according to David Yates. At OC-48 and up, however, SONET flows get their own wavelengths alongside the Ethernet *lambdas* (Figure 3).

### The Ghost Of Protocols Past

The diversity of Ethernet/TDM co-existence techniques should alert the 802.17 working group to proceed with caution. Equipment vendors are likely to ignore any standard that requires massive re-engineering to support TDM, especially since carriers are mostly happy to run their voice circuits over SONET.

Likewise, if technical complexity makes RPR too expensive, some other flavor of optical Ethernet will capture the market. In short, if the working group tries to be all things to all carriers, standards-based RPR products will never make it past the interoperability demos.

Networking history already has at least three defunct packet ring standards—FDDI, DQDB and token ring. And who even remembers failed convergence technologies like FDDI II and Iso-Ethernet? To avoid membership in this unfortunate club, the 802.17 working group should stick to its knitting—efficient, robust, metro-area data transport—and leave grander visions for another day□

**Companies Mentioned In This Article**

Appian Communications (www.appiancom.com)
Atrica (www.atrica.com)
Cisco (www.cisco.com)
Dynarc (www.dynarc.com)
IEEE 802.17 Working Group (www.ieee802.org/17/)
Lantern Communications (www.lanterncom.com)
RPR Alliance (www.rpralliance.com)