Resilient Packet Ring Alliance

RPR Business Case Study

Deploying RPR in a SONET/SDH Environment

May 2003

Contents

Copyright RPR Alliance 2003

Introduction

This case study demonstrates how Resilient Packet Ring (RPR) technology can augment the capability of SONET/SDH networks by enabling efficient and cost-effective transport of data services as well as traditional TDM traffic. The benefits for carriers of RPR become apparent as the network scales in terms of number of users as well as with increase in bandwidth per user.

Note: This is the first in a series of business cases developed by the RPR Alliance (http://www.RPRAlliance.org). The RPR Alliance strives to be the definitive technology resource for RPR. For a complete overview, see the RPR Alliance Technology Resource Center in the Whitepapers section of the RPR website - http://www.rpralliance.com/index.cfm?action=technology_white.

Network Description

The Starting Point: SONET/SDH

A typical SONET/SDH network from the customer premises to the point of presence (POP) is depicted below in Figure 1. In this network architecture, end users are connected to the network through T1 (DS1), T3 (DS3), OC3 interfaces, or through Ethernet interfaces (not shown in the picture). The hub node (DACS) switches both intra-ring as well as inter-ring traffic. For this example, we use an OC48 ring for metro access and an OC192 for the regional metro.

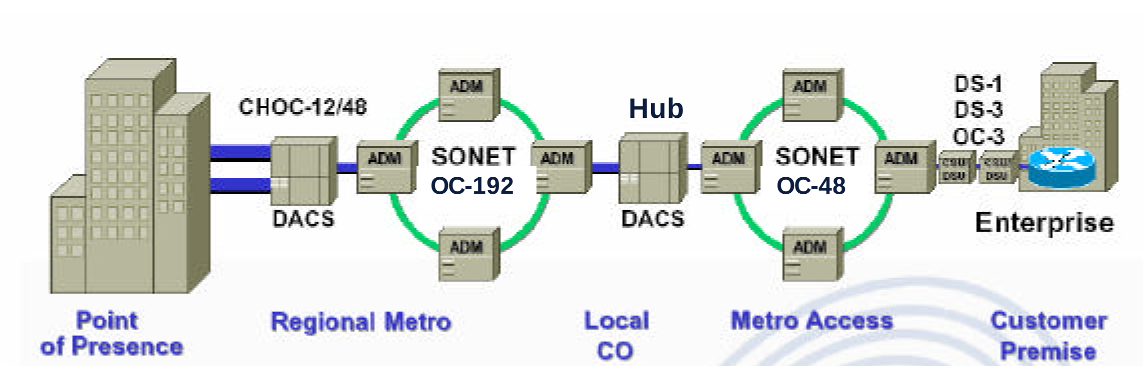

Figure 1: Legacy SONET/SDH, TDM-based Network

Copyright RPR Alliance 2003

Adding RPR Technology to the Network

To this existing infrastructure we introduce the use of RPR technology in Figure 2. In the RPR[1] network, similar to the SONET/SDH network, we have an OC48 ring for metro access and an OC192 ring for regional metro. In contrast, however, the customer interfaces are 10/100 Ethernet and 1 GbE. To offer Ethernet services to end users, we assume that fiber is available in the last mile from the central office (CO) to the customer premise. Also the hub node primarily switches traffic between rings and not within a ring as in SONET/SDH.

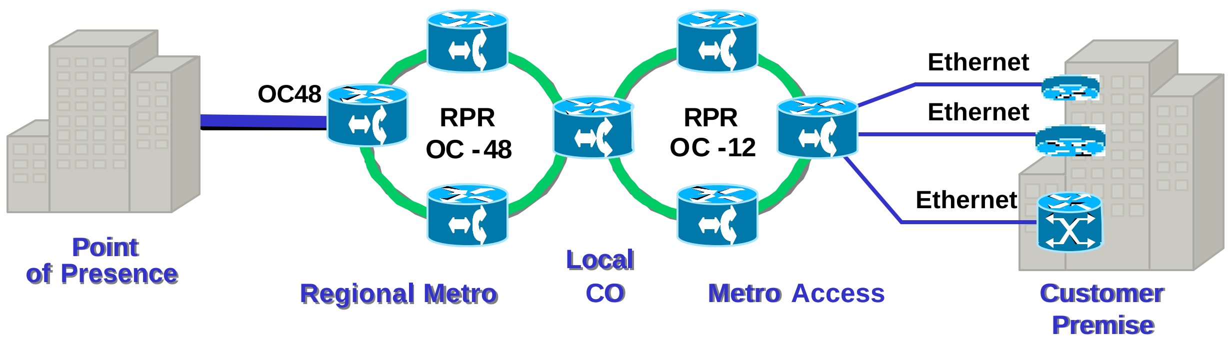

Figure 2: Typical RPR Network

Business Case Study

We have constructed a realistic traffic and subscriber growth profile that we describe in five scenarios to compare RPR CapEx and OpEx to that of traditional SONET/SDH. We have assumed, as our starting point, a network of 100 users evenly distributed on five nodes (SONET/SDH or RPR ADMs). We are not as concerned with actual costs as these costs are subject to many variables, and hard data are difficult to find. Therefore, for the sake of simplicity, even though RPR networks can be shown to be less expensive for the initial network build-out, we assume similar CapEx for both SONET/SDH and RPR networks in the study. This allows us to observe how CapEx and OpEx change as we scale the network.

---

[1] RPR networks can be deployed in greenfield networks or in existing SONET/SDH networks without a forklift upgrade either to the network or to the equipment.

Copyright RPR Alliance 2003

Scenario 1: Initial Deployment

This baseline scenario shows an initial deployment of 1.5 Mbps service to 100 subscribers. End users access this bandwidth through T1 connections in the case of SONET/SDH and through 10/100 Mbps Ethernet interfaces in the case of RPR. (Scenario 5 analyzes the case where Ethernet interfaces are available on the SONET/SDH network.)

If we take a current "average" price for both T1 and 10/100 Ethernet service to be around $1000/month,[2] this initial deployment scenario shows a comparable cost structure for both implementations. Both service providers and end-users are likely to see little difference in costs at this level of customer uptake. In fact, it is not unlikely that both RPR-based services and traditional TDM services will co-exist harmoniously in a real-world network, so it is good to know they have similar price-points from the early days of a new service rollout.

Scenario 1: Initial deployment with 1.5 Mbps service per subscriber		
	TDM only	**RPR**
Number of nodes (SONET/SDH or RPR ADMs)	5	5
Subscribers per node	20	20
Total subscribers	100	100
Bandwidth per subscriber	1.5Mbps	1.5Mbps
Bandwidth required per node	20x1.5Mbps=30Mbps	20x1.5Mbps=30Mbps
Total bandwidth required	5x30Mbps = 150Mbps	5x30Mbps = 150Mbps
Access ring speed	OC-48 (2.4 Gbps)	OC-48 (2.4 Gbps)
Effective ring bandwidth	2.4Gbps	4.8Gbps
Port count	100 T1 ports	100 10/100 Ethernet ports
Per port cost/month	$1,000	$1,000
Number of access rings required	1	1

Scenario 2: Increase the Per-User Bandwidth in a Small Increment

In this scenario, the bandwidth per user increases slightly to 5 Mbps from 1.5 Mbps. To accommodate this increase in the legacy SONET/SDH example, we need fractional T3s (6 Mbps). The customer premise equipment with T1s has to be upgraded to T3s, which involves truck rolls as well as costly hardware upgrades to ADMs with higher speed cards. Furthermore, the service provider has to tear down the T1 circuits previously established and provision fractional T3s all the way from customer premise equipment to each ADM on the network. This simple upgrade from T1 to fractional T3 increases both CapEx in terms of hardware costs and OpEx for truck rolls and provisioning costs.

---

[2] See the Appendix for pricing sources for T1, fractional T3s, 10/100Mbps Ethernet, 1GbE, etc.

Copyright RPR Alliance 2003

With RPR, the initial 10/100 port easily accommodates the increased bandwidth. Hence no truck roll, costly upgrades to the network, or bandwidth provisioning is needed.

Scenario 2: Increase the BW/user to 5 Mbps with fractional T3 interfaces		
	TDM only	**RPR**
Number of nodes (SONET/SDH or RPR ADMs)	5	5
Subscribers per node	20	20
Total subscribers	100	100
Bandwidth per subscriber	6Mbps (Frac T3)	5Mbps
Bandwidth required per node	20x6Mbps=120Mbps	100Mbps
Total bandwidth required	5x120Mbps = 600Mbps	5x100Mbps = 500Mbps
Access ring speed	OC-48 (2.4 Gbps)	OC-48 (2.4 Gbps)
Effective ring bandwidth	2.4Gbps	4.8Gbps
Port count	100 fractional T3 (6Mbps) ports	100 10/100 Ethernet ports
Per port cost/month	$2,600	$1,000
Number of access rings required	1	1

Scenario 3: Increase the Per-User Bandwidth Rapidly

We now portray a real-world scenario whereby 80 of our subscribers wish to upgrade their service from 5 Mbps to 40 Mbps, 10 subscribers to 100 Mbps, and the rest to 1 Gbps. In the case of legacy SONET/SDH, this upgrade of service would result in a complete ring redesign. Existing TDM-mapped cards must be replaced with new higher-speed hardware (T3, OC3, and OC48) and the operational expense substantially increases as technicians must travel to reconnect customers with higher speed interfaces. There will also be increased capital expenditures on additional access rings, because the first ring will run out of free bandwidth. In this example, we need a total of 13 rings. This, in turn, will require additional gear on the regional metro ring as well as a manual reconfiguration to setup new channels, both of which increase operational expenses.

In comparison, the RPR-based solution will scale service without any hardware or architectural changes. No architectural change is required in RPR as it supports strong over-booking of traffic due to statistical multiplexing around the ring and through spatial reuse. The 1 Gbps service can be provisioned through relatively inexpensive GbE ports. However, a truck roll is necessary to install these GbE ports at the customer premises.

Copyright RPR Alliance 2003

Scenario 3: Rapid increase in per user bandwidth		
	TDM only	**RPR**
Number of nodes (SONET/SDH or RPR ADMs)	5	5
Subscribers per node	20	20
Total subscribers	100	1000
Each Node composition:		
Number of subscribers receiving 40 Mbps (T3)	16	16
Number of subscribers receiving 100 Mbps (OC3)	2	2
Number of subscribers receiving 1 Gbps (OC48)	2	2
Bandwidth required per node	16x45Mbps+2x155Mbps+ 2x2.4Gbps=5.83Gbps	16x40Mbps+2x100Mbps +2x1Gbps=2.84Gbps
Total bandwidth required	5x5.83Gbps=29.15Gbps	5x2.84Gbps=14.2Gbps
Oversubscription	"1:1"	"1:3"
Access ring speed	OC-48 (2.4 Gbps)	OC-48 (2.4 Gbps)
Effective ring bandwidth	2.4Gbps	4.8Gbps
Port count	80 T3 10 OC3 10 OC48	90 10/100 Ethernet ports 10 1GbE ports
Per port cost/month	T3: $25,000 OC3: $50,000 OC48: $450,000	10/100: $1000 1GbE: $10,000
Number of access rings required	29.15Gbps/2.4Gbps=13	14.2Gbps/(3x4.8Gbps)=3

Scenario 4: Increase the Number of Users by Tenfold

In the previous example we stressed the network by increasing the bandwidth per user. As we evolve the network, we now want to study the effect of increasing the number of users by tenfold while keeping the same bandwidth requirement mix as in Scenario 3 (80%a t 40 Mbps, 10% at 100 Mbps, and 10% at 1 Gbps services).

The TDM-only example results in a substantial increase in CapEx and OpEx to deploy and maintain 122 rings. While in RPR example, the CapEx and OpEx increases are relatively lower as we need to build only 10 rings.

Copyright RPR Alliance 2003

Scenario 4: Increase the number of users by tenfold		
	TDM only	**RPR**
Number of nodes (SONET/SDH or RPR ADMs)	50	50
Subscribers per node	20	20
Total subscribers	1000	1000
Each Node composition:		
Number of subscribers receiving 40 Mbps (T3)	16	16
Number of subscribers receiving 100 Mbps (OC3)	2	2
Number of subscribers receiving 1 Gbps (OC48)	2	2
Bandwidth required per node	16x45Mbps+2x155Mbps +2x2.4Gbps=5.83Gps	16x40Mbps+2x100Mbps+ 2x1Gbps=2.84Gbps
Total bandwidth required	50x5.83Gbps=291.5Gbps	50x2.84Gbps=142Gbps
Oversubscription	"1:1"	"1:3"
Access ring speed	OC-48 (2.4 Gbps)	OC-48 (2.4 Gbps)
Effective ring bandwidth	2.4Gbps	4.8Gbps
Port count	800 T3 100 OC3 100 OC48	900 10/100 Ethernet ports 100 1GbE ports
Per port cost/month	T3: $25,000 OC3: $50,000 OC48: $450,000	10/100: $1000 1GbE: $10,000
Number of access rings required	291.5Gbps/2.4Gbps=122	142Gbps/4.8Gbps=10

Scenario 5: Compare RPR with Next-Gen SONET/SDH (VC, Ethernet Ifs)

Much of the scalability issues seen in the legacy network have to do with the traditional TDM bandwidth granularity available in the SONET/SDH environment. The resultant "stranded bandwidth" effect has been well documented and is being addressed by next generation solutions that significantly address this problem. To account for these improvements, we propose Scenario 5. Here we use the same number of subscribers and bandwidth mix as in the Scenario 4 but deploy native Ethernet subscriber interfaces for the SONET/SDH network. We also take advantage of virtual concatenation (VCAT) technology to efficiently transport data over SONET/SDH. As an example, with VCAT, fast Ethernet (100 Mbps) can be transported in two discrete STS1s (~2x52=104 Mbps) rather than in an STS3 (155Mbps) as before. Similarly, one gigabit Ethernet will be transported using 21 discrete STS1s rather than with 48 concatenated STS1s.

As we notice from the following table, while the aggregate SONET/SDH bandwidth requirement drops considerably (due to VCAT), we still need 65 rings to provision that bandwidth. The number of rings continues to be high in SONET/SDH as it does not support oversubscription, which is a key RPR benefit.

Copyright RPR Alliance 2003

Scenario 5: Same as scenario 4 but with VC based SONET/SDH network		
	TDM only	**RPR**
Number of nodes (SONET/SDH or RPR ADMs)	50	50
Subscribers per node	20	20
Total subscribers	1000	1000
Each Node composition:		
Number of subscribers receiving 40 Mbps (T3)	16	16
Number of subscribers receiving 100 Mbps (OC3)	2	2
Number of subscribers receiving 1 Gbps (OC48)	2	2
Bandwidth required per node	16x45Mbps+2x104Mbps +2x1.092Gbps=3.112Gbps	16x40Mbps+2x100Mbps +2x1Gbps=2.84Gbps
Total bandwidth required	50x3.112Gbps=155.6Gbps	50x2.84Gbps=142Gbps
Oversubscription	"1:1"	"1:3"
Access ring speed	OC-48 (2.4 Gbps)	OC-48 (2.4 Gbps)
Effective ring bandwidth	2.4Gbps	4.8Gbps
Port count	900 10/100 Ethernet ports 100 1GbE ports	900 10/100 Ethernet ports 100 1GbE ports
Per port cost/month	10/100: $1000 1GbE: $10,000	10/100: $1000 1GbE: $10,000
Number of access rings required	155.6Gbps/2.4Gbps=65	142Gbps/(3x4.8Gbps)=10

RPR Business Benefits

As we've seen, a key benefit to the RPR deployment lies in its scalability. As the network grows to support increased demand in terms of subscriber uptake and bandwidth increases, carriers can realize compelling benefits. In this section, we will delve into certain RPR features that make these business benefits possible.

Bandwidth Efficiency

RPR offers Ethernet-like bandwidth efficiencies while preserving the carrier-class features such as 50-ms protection and quality of service for voice. This is possible through following features:

- Dual-ring architecture for working traffic
- Spatial reuse
- Statistical multiplexing
- Efficient multicast support

Each one of these features described below in more detail offers compelling benefits for carriers to reduce CapEx and OpEx as well as increase revenue-generating traffic in their networks.

Copyright RPR Alliance 2003

Dual Ring Architecture

RPR can use both rings to carry working traffic, immediately resulting in twice the amount of bandwidth compared with traditional SONET/SDH where half the bandwidth is allocated for protection traffic. This can be used to advantage in the case where mixed service levels are carried on the ring as lower classes of service can yield to higher ones in the case of a fault. In case of a fiber cut or a node failure, the RPR protection mechanism will steer or wrap traffic away from the fault within 50ms, thereby maintaining the quality of service subscribers have come to expect in MAN services. The RPR fairness algorithm will ensure that guaranteed classes of traffic continue to be serviced according to their individual service level agreements.

Spatial Reuse

Unlike in certain SONET/SDH architectures (for example, UPSR), RPR does not lock out the entire ring for the duration of a data transfer. Bandwidth is only "consumed" on the segments between source and destination nodes. This allowing other transfers to take place simultaneously between other nodes on the ring as shown in Figure 3 This feature, known as spatial reuse or bandwidth multiplication, significantly increases the efficiency of the ring operation.

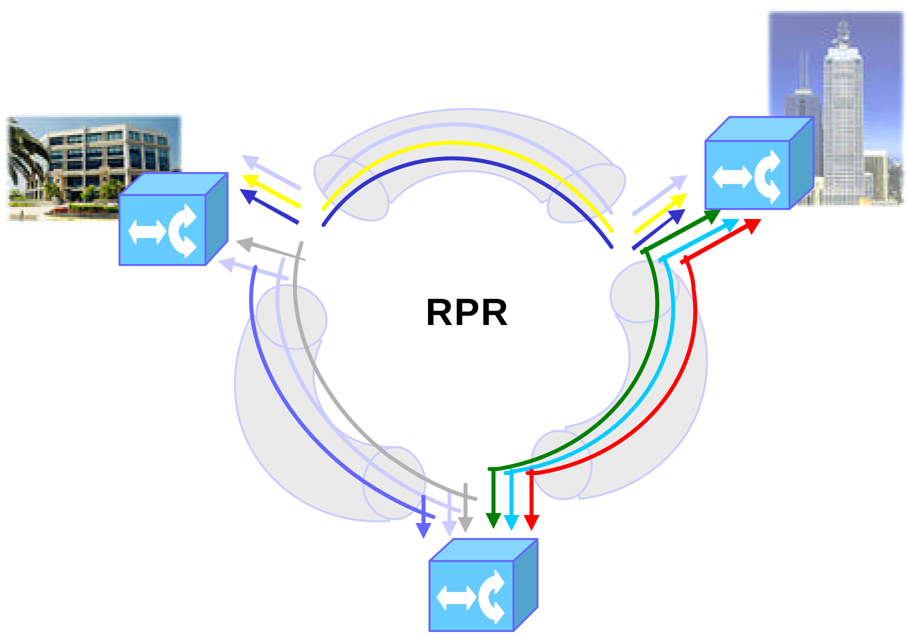

Figure 3: Spatial Reuse in RPR Networks

Copyright RPR Alliance 2003

Statistical Multiplexing

RPR supports a very efficient bandwidth sharing capability through ring-wide statistical multiplexing. This resembles a distributed switch around the ring resulting in greatly reduced back-hauling of traffic within the ring. In RPR, a hub node is only required to backhaul inter-metro traffic to the backbone network but not for intra-metro switching of traffic. A true any-to-any connectivity is thus possible around the MAN. In contrast, circuit-oriented alternatives require multiple point-to-point connections (each individually provisioned and managed, and any bandwidth provisioned can no longer be shared). This results in more traffic that has to be backhauled to a centralized hub for switching. In addition, to improve the performance of statistically multiplexed networks, the ability to handle bursty data traffic is necessary. RPR shares idle bandwidth between all stations on the ring. RPR bursts are limited only by the capacity of the customer interface (100Mb/second or 1Gb/second for Ethernet) and the total capacity of the ring (1-10Gb/second). Figure 4 shows how RPR can distribute spare bandwidth to accommodate burtsy data sources.

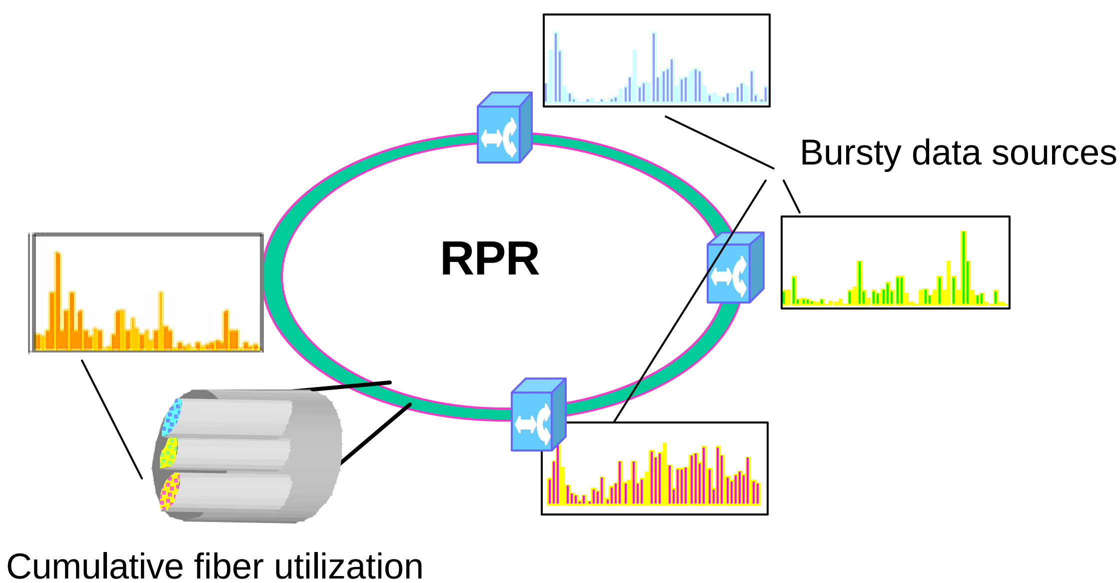

Figure 4: Statistical Multiplexing in RPR Networks

Copyright RPR Alliance 2003

Multicast

Due to the nailed-up nature of a circuit-oriented network, multicast traffic wastes bandwidth, especially when bandwidth-intensive multi-point applications such as distance learning and video streaming are considered. As shown in Figure 5, multiple copies of multicast packets must travel around ring (one copy per circuit), wasting considerable bandwidth. In the case of the packet-aware ring enabled by RPR, a single packet copy is sent around the ring and copied at each valid destination, thereby greatly reducing the bandwidth requirement on the ring for multi-point applications.

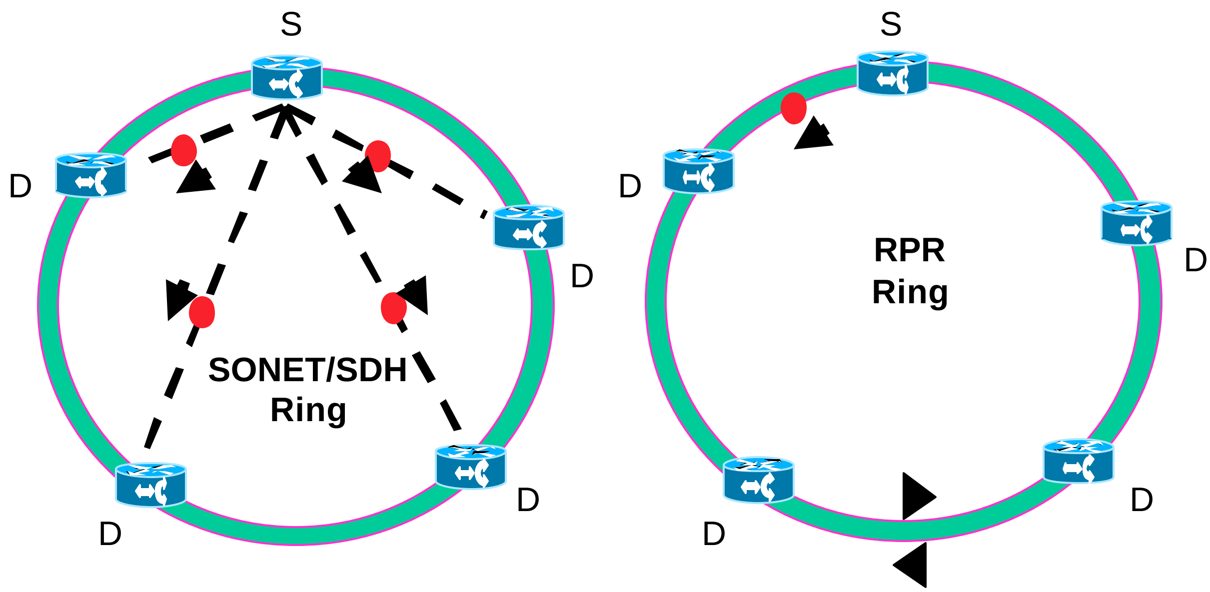

Figure 5: Multicast in RPR Networks

Port Consolidation

A typical point of presence employing a traditional point-to-point architecture requires numerous expensive line interfaces on the core router as shown in Figure 6. In contrast, the RPR network architecture elegantly consolidates these port connections into interfaces to the ring itself. User experience has shown up to a 50% cost savings in CapEx. OpEx is greatly improved as network management and sparing requirements are reduced by RPR.

Copyright RPR Alliance 2003

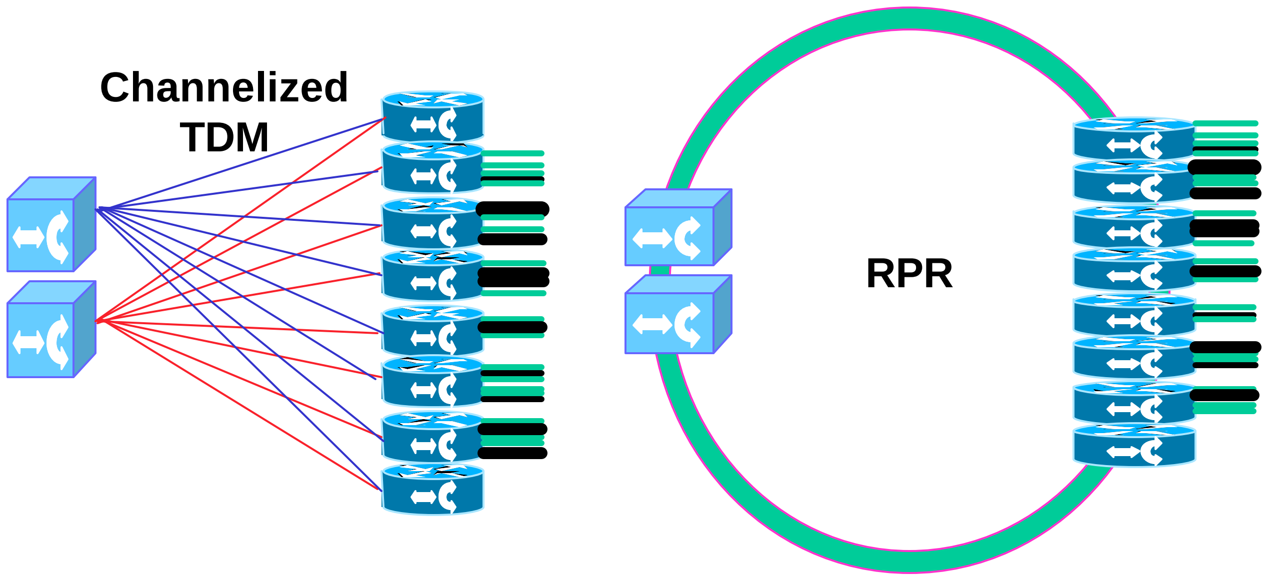

Figure 6: Port Consolidation in RPR Networks

Simplified Service Provisioning

A common complaint from data service customers is that it takes too long for carriers to provision services. Activation times on the order of six weeks to six months for DS1 and DS3 services are quite common, with services at OC-3 rates and higher taking even longer.

A significant portion of this delay can be attributed to the underlying TDM infrastructure and its circuit-based provisioning model. Creating an end-to-end circuit takes many steps. First, the network operator identifies the circuit's physical endpoints to the management system. The operator must then configure each node within the ring for all the required pass-through and add-drop connections, some of which might require re-engineering or additional hardware. This provisioning operation is time- and labor-intensive.

Newer transport systems automate some of the provisioning steps, but the network operator still needs to perform traffic engineering manually to optimize bandwidth utilization on the ring. The operator must be aware of the network topology, the traffic distribution on the ring, and the available bandwidth on every link traversed by the circuit.

Copyright RPR Alliance 2003

An RPR system, by contrast, offers a very simplified service model. The ring functions as a shared medium where all stations on the ring share bandwidth without predefined circuit restrictions. Each station has visibility into the current bandwidth capacity available on the ring and can therefore provision new services far more simply. There is no need for a node-by-node and link-by-link capacity planning, engineering, and provisioning exercise. The network operator simply identifies a traffic flow and specifies the quality of service and bandwidth a customer should receive on the ring. Limiting the provisioning tasks to the customer interfaces significantly improves operational ease and the economics for carriers.

Circuit Emulation

While most of our analysis has centered on data services, it is important to realize that most carrier revenue today comes from voice services carried over the legacy TDM network. We have explored how RPR can be deployed in parallel with such legacy services affording incremental revenue opportunities without impacting current operational procedures. As the network evolves, circuit emulation of TDM services over RPR further enables carriers to keep their current high revenue TDM services while migrating them to the RPR portion of the network. The circuit emulation services are made possible by the reduced latency and jitter of RPR as well as its ability to deliver multiple service qualities on a common, shared infrastructure.

Conclusion

We have shown how RPR facilitates the scaling of an existing network to support increased traffic and subscriber levels in a metro environment. Leveraging these benefits, carriers can significantly reduce their CapEx and OpEx in addition to increasing their revenue generating traffic with a network that can scale for their future needs.

About the RPR Alliance

The RPR Alliance is an industry advocacy group committed to the development of an RPR technology standard for the networking industry. The Alliance promotes the adoption of an RPR standard for LANs, MANs, and WANs by educating the networking industry about RPR technology and the benefits of an IEEE standard as well as by fostering multi-vendor interoperability. Principal members of the RPR Alliance include Cisco Systems, Corrigent Systems, and Nortel Networks. Alliance Semiconductor, Cortina Systems, Infineon Technologies, Intel, Ixia, and Lantern Communications are participating members in the Alliance. For more information about the Alliance and the membership application, see http://www.RPRAlliance.org.

Copyright RPR Alliance 2003

Appendix

Service Pricing Details

Service	Price/month	Source
Full T1	$1000	www.t1solutions.com
T3 (6Mbps)	$2600	www.coop.net
Full T3	$25,000	www.bandwidthsaving.com
OC3	$50,000	www.bandwidthsaving.com
OC48	$450,000	www.ecttelecom.com
10/100 Mbps Ethernet	$1000	www.cogentco.com
1000 Mbps Ethernet	$10,000	www.cogentco.com

These prices are valid as of March 2003.

Note: The RPR Alliance has worked diligently to develop realistic scenarios and their associated costs based on readily available data from vendors in order to provide a reasonable comparison. Actual costs may vary. Considerable effort has also been made to compare aspects of the network (such as the bandwidth or number or rings consumed) that allows a fair comparison without relying on these discrete costs.

Copyright RPR Alliance 2003