Technical Paper
Cable-Tec Expo® 2002
Topic Category: Deploying VOD that scales

VOD Everywhere! Considerations in Transport Methods for Scalable VOD/SVOD Deployment

Jason Shreeram, Product Manager, Prisma Optical Networks
Scientific-Atlanta, Inc.

Abstract

The new "Everything-on-Demand" (EOD) model is challenging Broadband Hybrid-Fiber Coax network engineers to construct the systems today that will allow offerings in new services, compete in non-traditional and yet-to-be-defined markets, and increase value for MSO shareholders. As operators introduce these new interactive services, they must adapt existing system infrastructures to accommodate the additional equipment that an EOD system requires. The necessity of building these networks once again challenges the industry to stack yet another network on top of existing analog video, digital video, CMTS, and business data networks, or develop a migration strategy to ensure future scalability in the network as services develop.

Deploying video-on-demand (VOD) and other high-bandwidth interactive services requires solving many technical and logistics challenges. Fortunately, recent innovations and practices make dealing with these challenges easier by reducing the burden on new headends that must deploy on-demand technologies. For example, Gigabit Ethernet (GbE) systems have recently been developed to allow VOD streams to be transported throughout the network at a very high capacity in a cost-efficient manner. The relevant options and considerations for these GbE technologies and products used with on-demand applications are discussed in this paper. We will discuss both "stacked" and "converged" network architectures for transport and delivery of VOD services and the key considerations for determining which methodology to use.

The EOD Model

Everything-on-Demand (EOD) is a fairly new acronym that provides a catch-all phrase to include all video content transmissions and exclude none. As late as last year, the models for Hybrid-Fiber Coax networks were being developed with the idea that the Cable industry would be deploying Movies-On-Demand (MOD) in the near future. However, the successful trials of Subscription Video-on-Demand (SVOD) have forced us into rethinking the architectures and reassessing the statistical models we thought were valid last year. For example, a typical usage analysis of a MOD model is extremely “bursty” with regard to its traffic pattern, yet highly predictable in resource requirement (Figure 1):

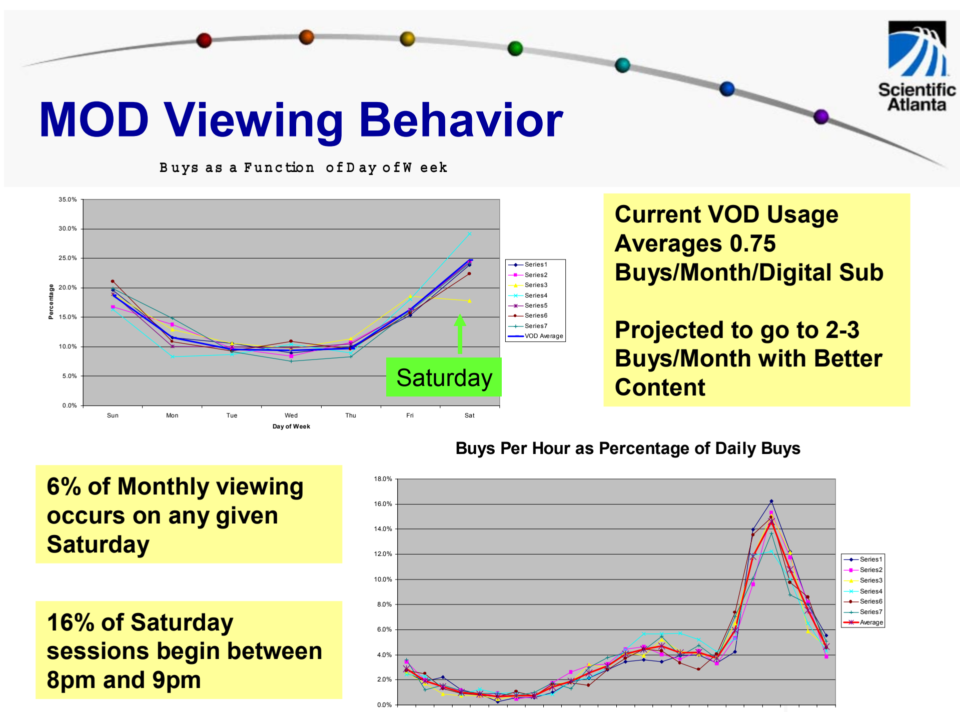

Figure 1 MOD Viewing Behavior

Current trend analysis shows that the MOD model is extremely time predictable and, as shown, this can help with trafficking requirements. During peak viewing times (Saturday nights between 8 and 9PM) the traffic spike can help the operator to determine the bandwidth requirement for any given site and the statistical probability of any blocked or “busy” responses. According to the statistical evidence accumulated thus far, approximately 6% of the monthly viewing of any population occurs on any given Saturday night (Wall, CT 2001). Presently the MOD model is limited by the content of the “On-Demand” titles. In June 2001, when this data was accumulated, the average of .75 MOD buys per month per digital subscriber was limited by the title content available at the time. It is expected that when the content agreements can provide first run movie titles on par with that of the video rental stores, this average of .75 will increase to between two to three buys per month per digital subscriber. An effective network design will manage this “peak flow” effect and provide for an acceptable system response level. The operator must determine what the acceptable number, if any, of “no service available” is and the probability of losing that customer business to a competitor such as Video rental stores or DBS subscriptions. This growth and service level will determine the over-provisioning of the network required to properly allot traffic to the network.

The Subscription Video-on-Demand (SVOD) model has very little similarity to the aforementioned MOD model. SVOD subscribers have much faster session set-up and tear down turnaround than MOD viewers, primarily because the nature of the programming is much different. In this case, the average SVOD viewer uses the system much more often, averaging between .6 and .8 sessions per day, and the viewing trend analysis shows much flatter behavior, with less traffic peak at the “prime time” Saturday night viewing time than the MOD model. In fact, as Figure 2 illustrates, the standard user rates look more like data from internet and business models than video on demand models. The peak sessions that did occur contained <5% of the total SVOD subscribers at any given time.

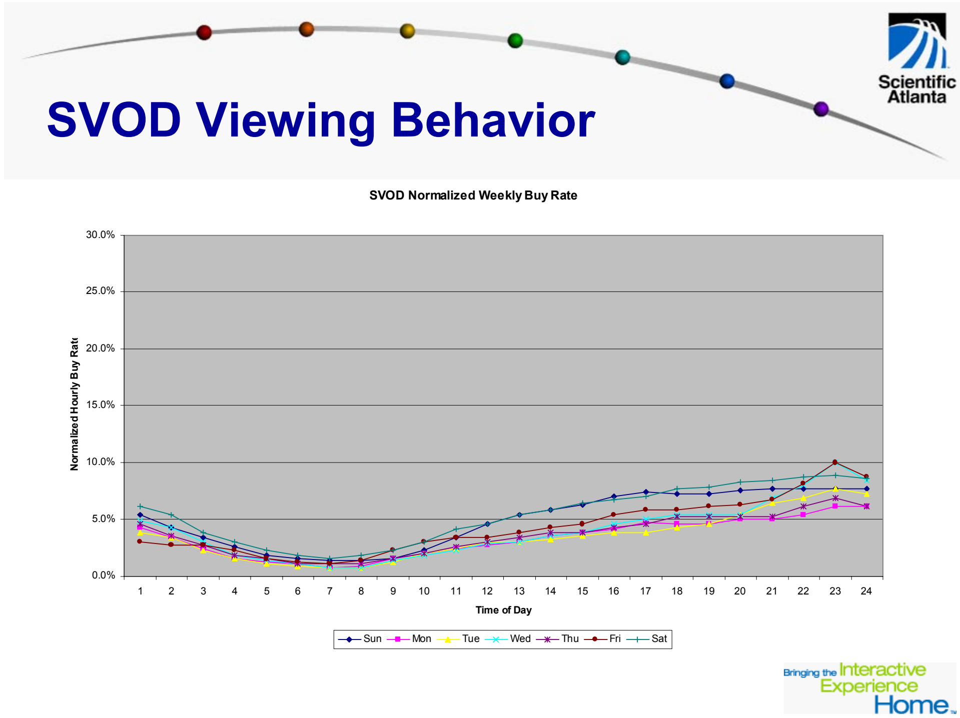

Figure 2 SVOD Traffic Rates

Since the trend of a SVOD model is similar to the idea of an EOD model, we can then logically scale the requirements of our network capacity to model future traffic requirements and extrapolate that data to our network backbones. The idea that streaming video will originate anywhere in the IP cloud and aggregate at a single point of network entry is not new and is being realized today. Many IP video sources are becoming available as high speed internet connections roll out, and the storage of video content is rapidly becoming big business.

VOD Network Developments

The newer SVOD and EOD models, coupled with advances in server technologies and IP migration into HFC networks, have created significant new opportunities for transport considerations. Whereas last year most of the networks being deployed used either proprietary ASI digital transport or analog DWDM systems, changes in both edge devices and system architectures have occurred to allow operators more choices in transport. Presently there are three common methods that operators are using to transport VOD streams from primary headends to hubs: Distributed servers; Centralized server/distributed QAM's; and Fully Centralized servers and QAM's. The criteria for determining transport types are pre-determined by the selection of either a centralized server or distributed server architecture (Figure 3).

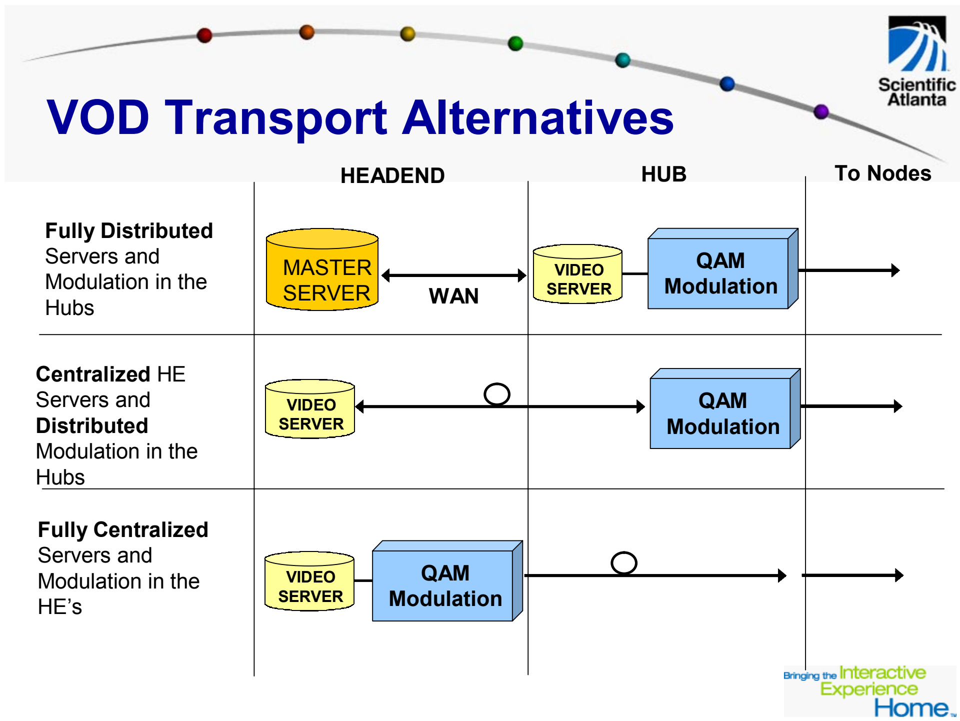

Figure 3 VOD Transport Alternatives

Distributed Architectures

Distributed architectures typically consist of a master headend system that provides control of the VOD file servers placed at multiple primary or secondary hub locations. These hub-based VOD servers store and stream the video to the serving areas and ultimately to the digital set-top converter at the home. Hybrids or variations of these architectures do occur. An example of these might be where a "master library" of all programs is kept at the master headend location to serve the whole network as well as closely surrounding areas. Additionally, each hub contains servers with the frequently-used program content (such as new releases), and these servers provide for the immediate or surrounding areas. The main library server "refreshes" the hub servers as new content is required or on a predetermined basis. Most distributed architectures will contain some of these aspects, where the hub servers are refreshed with new program content from a master server library or where a larger server at the master headend has an expanded library of programs that can be accessed.

Centralized Architectures

Centralized architectures place all the VOD servers at the master headend location. From the master headend, the VOD streams are transported to the primary and secondary hubs (if any), and are then sent to the nodal area and ultimately to the digital set-top box at the customer premises. Centralized architectures have two variations that are predominant in today's architectures.

The first variation can be characterized as "fully centralized," meaning both the servers and the QAM modulators are co-located together in the headend. The transport of QAM modulated digital signals is accomplished via analog lasers on DWDM backbones. This equipment transmits the signals around a fiber optic ring to the primary hub and/or secondary hub locations. At the hub, the DWDM multiplexers and receivers separate the signals before electrically combining them with the broadcast signals for distribution using standard 1310nm or 1550nm optics, amplifiers, and passive optical networks. DWDM is normally used to minimize the number of fibers required between the master headend and primary hub/secondary hub and for expansion capability. Analog DWDM deployments are prone to fiber dispersion limitations and usually are used in systems where the headends and hubs are within 50 to 60 km of each other.

The second centralized architecture can be considered a "centralized server; distributed QAM" topology that utilizes various forms of digital laser equipment to distribute baseband digital signals on DWDM backbones. Servers today utilize digital transmission formats such as Asynchronous Serial Interfaces (ASI) and Gigabit Ethernet (GbE). These products utilize direct feeds from the media server outputs and transmit those feeds directly onto a wavelength in the DWDM ITU grid. The efficiency of using the direct feeds and transmitting in the digital spectrum is impressive. Not only are there technical advantages of digital transmission pertaining to fiber dispersion effects, but there are significant cost advantages as well. Additionally, most of the major server vendors have announced plans to migrate to a GbE interface in the 2002 timeframe, so this format looks to be a logical choice for transmission media.

Stacked Networks Utilizing ASI and GbE Backbones

Considerations for VOD deployments can be challenging in today's environment. The deployment of stacked networks has occurred as a phenomenon rather than a plan. As applications have developed that require transport, the operator adds the equipment, finds the capacity, dedicates a fiber or two, and deploys. In the late 1980s and early 1990s, operators started with basic fiber backbones and, over time, applications such as CMTS, digital video, HITS, and switched circuit telephony have all added equipment and capacity requirement to the network. In some cases, operators have had the good fortune of possessing adequate facility and fiber growth capacity to roll these services out without much hesitation, but for others, buildings have been rebuilt, fibers have been overlaid, and new plant has been constructed. This stacked network has led to a "silo effect" in the headends, whereby four different sets of equipment are run by four different groups of engineers, many of whom share the same fiber (Figure 4).

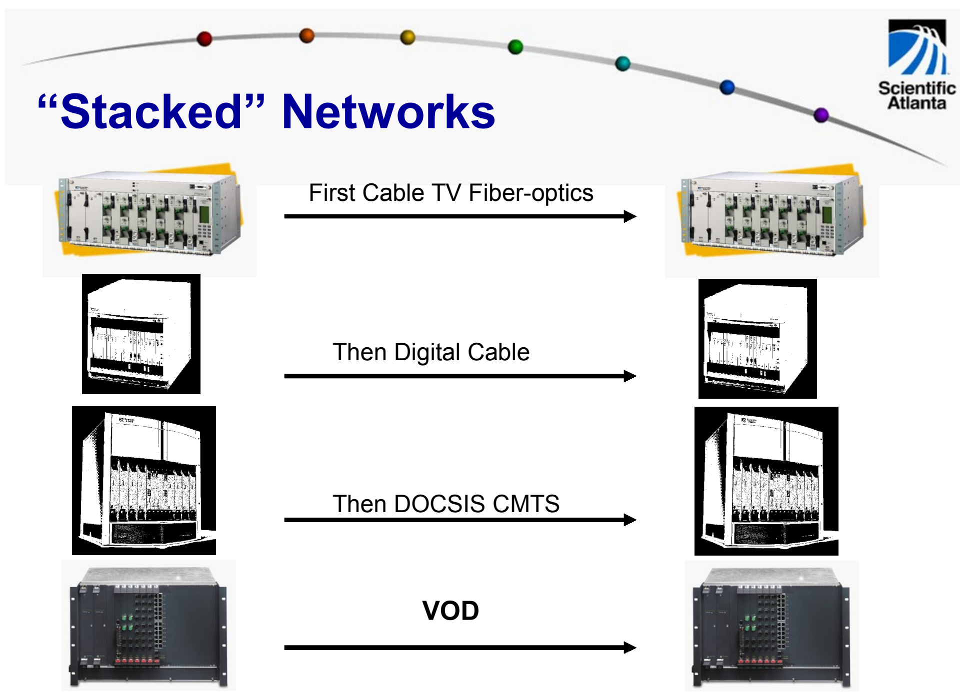

Figure 4 Stacked Networks

Over time, these stacked networks may have challenged the fiber capacity of the networks, which required the addition of Dense Wave Division Multiplexing (DWDM) equipment to enhance transmission capacity. The addition of EOD will challenge the networks once again, as will rollouts of new services such as Voice over IP (VoIP) and DOCSIS1.1. Once again operators are faced with looking short term vs. long term. In today's VOD environment, the launching of ASI proprietary transport is one such short-term solution. There is no migration path from an ASI transport to any of the future rollouts mentioned above. However, there is some relief coming in the form of GbE transport. Server vendors and QAM modulator manufacturers are now migrating product to a GbE standard which will allow operators the flexibility of choosing to design a network that can be scaled as IP services launch. The current format of server output which is limited to ASI puts the system architect in the position of choosing technology dedicated to the specific application of VOD. As we discussed earlier, the model of EOD predicts massive growth in the number of streams to be transported and the network capacity. By deploying VOD in an ASI output format, the nature of the system connection is a point-to-point network because, unlike Gigabit Ethernet, ASI cannot be switched. Consequently, the systems that deploy ASI transport have very little migration path in the electrical domain, and future growth scalability will be limited to fiber capacity rather than transmission rate. Additionally, the QAM modulators that are available in today's markets have either ASI or GbE inputs, therefore there are many variants of the networks being deployed today (Figure 5). Some of these variants might employ GbE backbone with ASI conversion in the QAM modulator or in a translation device placed in front of the QAM. This is an effective interim step to get networks rolled out today, but adds unnecessary cost to the transport system. The general migration to an all GbE network, however, has tremendous momentum and should be completed by the end of 2002.

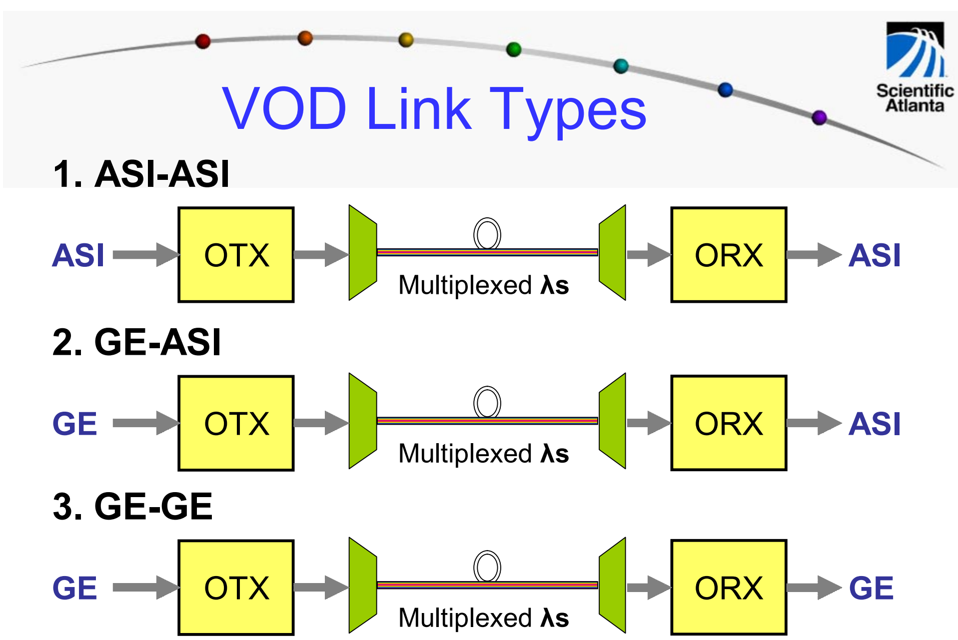

Figure 5 Variant Protocols

When considering system capacity and data rate throughput, most system operators prefer to use GbE transport because of the ability to both multiplex and switch the GbE transport, thus effectively utilizing bandwidth upon demand. Additionally, since GbE can provide “Jumbo packets” at 900Mbps, the transmission capacity is much higher per fiber than multiplexed 270 Mbps ASI transport. As described earlier, the transport requirements of SVOD transport are much less peak oriented than the MOD models. Therefore, the ability to switch the network capacity as required can effectively provision server capacity and partition system capacity to where it is needed. Also, the ability to multiplex numerous GbE inputs onto a high capacity laser can save on the cost of the lasers and improve the system’s capacity. The key consideration of the network architect should be to maintain effective throughput in the network, always insuring the limit of throughput capacity is regulated by the edge devices (either the server or the QAM modulator) rather than the network. If GbE switching is utilized, this can be accomplished easily with drop-in upgrades for scalability as traffic capacity requirements increase. In addition, the 10 Gigabit compliance testing schedule of the IEEE 802.3 has been completed and many vendors have launched transport at 10Gbps. The upgrade of VOD network backbones from 1 Gbps to 10 Gbps with full backward compatibility will lower transport cost, increase capacity, and provide the industry with a scalable and growth-oriented platform for the EOD models under consideration. As more IP type platforms develop and further IP traffic is required on the HFC network, this GbE scalability will enhance the value of the network and provide solutions for business migration. For example, VoIP services can be easily migrated onto a switched 10Gbps ring and run harmoniously with VOD traffic. However, GbE networks do not add the intelligence required to set up business-to-business services, VoIP, and other redundancy critical features, especially when quality of service guarantees are required.

Converged Networks

Resilient Packet Ring as an Everything-On-Demand Transport Solution

Limitation to GbE Quality of Service functionality and system convergence are also pushing the network operators to consider an alternative GbE transport protocol called Resilient Packet Ring (RPR). In the past, SONET/ATM technologies were used to transport data and separate networks were required for different types of traffic. Today we are seeing the emergence of a new packet-based standard for transporting interactive video, data, and new services over optical networks.

Both the IETF and IEEE (802.17 Working Group) are now focusing on this innovative and emerging transport solution.

A Resilient Packet Ring (RPR) is an optical packet transport ring that can interact with multiple nodes on the ring. The nodes are responsible for removing and placing packets on the ring addressed to them. RPR technology provides more efficiency of bandwidth than a voice-optimized network using SONET-based equipment. RPR is “packet optimized” to accommodate emerging packet traffic while supporting legacy TDM interfaces and video interfaces over a common platform. With RPR-based platforms, the operator can deliver multiple services on the network at the same time. Data and voice are combined into packets and are placed on rings for transmission; narrowcast signals can also be combined with the broadcast spectrum.

The optical access switching ability of RPR-based platforms makes it a perfect solution for services which require increased bandwidth such as video on demand. In addition, since RPR networks run IP packets, RPR can be an effective interface for conversion of ASI to GbE or vice versa. In fact, RPR networks can provide all the translation necessary to interface any type of server to any type of QAM modulator without interface devices. RPR-based platforms create packet-based optical networks with the quality of service (QoS) guarantees and dynamic bandwidth management that compare with ATM networks.

RPR is a hybrid of circuit and optical access switching offering guaranteed on-demand delivery while eliminating the need for reserved bandwidth. Purely packet-switched, RPR is unlike traditional or hybrid SONET Add-Drop Multiplexers (ADMs) that cannot offer the same efficiency, bandwidth, QoS, or robustness that are critical for services that require guaranteed end-to-end packet delivery. A RPR-based platform can be deployed as a switched ring and/or mesh topology. RPR provides carrier-class transport with all carrier grade features such as equipment level protection through redundant hardware and software, hot-swappable and hot-standby modules, fault tolerant hardware and software, and NEBS 3 compliance.

Quality of Service Guarantees

The packet-switching architecture provides extensive QoS support including policing, shaping, class-based queuing, and flexible scheduling based on strict priority at the point where data enters the network. This architecture allows the operator to use Ethernet interfaces as WAN interfaces rather than the T1, DS3, OC-3 POS, and OC-12 POS interfaces. TDM traffic carried on the RPR-based platform is guaranteed the same QoS as traffic on a pure TDM network. SONET traffic is guaranteed the same QoS as traffic on a SONET network through WDM protocol-independent transport. Each line card on the RPR-based platform is a full duplex local switch. Switching fabric modules for network-side ring switching reside on the switching controller card. This two-level switching scheme ensures that local switched traffic and network switched traffic do not collide.

Innovations for Reliability and Resilience

RPR offers several innovative methods to ensure that the "packet ring" is reliable and resilient to breaks in the network:

- **Automatic topology discovery**.
  RPR-based platforms can detect what other devices are on the network and where they are located.
- **Sub-50 millisecond protection switching.**
  The dual rings of RPR allow information to be quickly recovered if the data transmission fails on the primary ring. The backup ring assumes the responsibility of the primary ring in less than 50 milliseconds. With a packet header format suitable for robust end-to-end packet delivery, RPR strives for the shortest header possible to indicate the destination for the packet in order to leave ample room for the payload.
- **Distribution of Stratum clocking across the ring.**
  Distribution of stratum clocking across the ring is required to enable end-to-end transport of DS1s through the ring without frame slips. An RPR-based network is fully Stratum 3, with each node containing multiple internal Stratum 3 clocks. Fair bandwidth allocation across TCP/UDP flows in over-provisioned or best-effort traffic classes. RPR-based platforms police the data so that services requiring large amounts of bandwidth do not prevent other services from being transmitted.

Conclusion

EOD traffic, while still yet to be implemented, will provide operators with significant challenges for managing traffic requirements. Initial SVOD trials have shown significant variances between MOD traffic and SVOD traffic. These variances in traffic requirements will require flexible, scalable connections for ease of demand-mapped service allocation. One method of providing this pathway to scalability is switched Gigabit Ethernet, which will provide a common IP traffic platform for both on-demand services and future platform enhancements such as Voice over IP. An ideal network for VOD services will allow interfaces for both ASI and GbE to be common from transmit to receive and allow for internal translation without adding unnecessary costs to the transport system. The advent of new technologies is forcing operators to once again determine if stacking yet another network on top of existing infrastructure is the best method, or if alternatives should be considered. The ideal platform is one that will provide scalability, connections from numerous interfaces, and IP functionality for future service launches. RPR can provide a solution to converge the network over time with seamless migration from ASI devices to GbE devices, yet provide an effective cost per stream VOD transport for today's on-demand services.

The author would like to acknowledge the following contributors from Scientific-Atlanta to this paper: Lorenzo Bombelli, Joachim Vanhaecke and Bill Wall.