PACKET NETWORK TOPOLOGIES FOR NEXT GENERATION VIDEO ON DEMAND AND SWITCHED BROADCAST SERVICE DELIVERY

John Amaral and Paul Pilotte
Artel Video Systems

Abstract

This paper focuses on packet network architectures that are optimized for the delivery of next generation Video-on-Demand and Switched Broadcast. The paper explores the behavior of switched video delivery networks that satisfy the growing user demands for unique orthogonal sessions. A detailed analysis of video delivery infrastructure composition is undertaken. The paper discusses packet switching systems, optical transport, Layer 2 forwarding, QAM modulation and storage. A hypothetical 300,000-subscriber VoD network is employed as the basis for describing network behavior under several scenarios. The analysis culminates in a cost-effective, extremely high capacity network that dramatically increases bandwidth resource utilization and provides dynamic and agile program delivery. The disclosed topologies possess effective redundancy and resiliency. Several practical examples are considered with regard to the disclosed topologies; the examples include "Everything on Demand" (EoD) and Switched Broadcast Services. The analysis is predicated by the feasibility and practicality of the described topologies. Considerations such as interoperability, cost, ability to deploy, and ease of use is taken into account as important factors when describing the topologies.

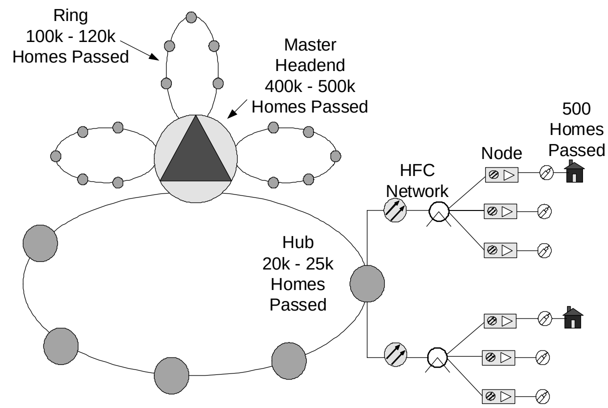

Figure 1 - Typical Cable Network

THE MOVE TO VOD

The deployment of next generation video on demand, EoD, and diverse content offerings by the MSO's in one form or another is regarded as a foregone conclusion. MSO's must deploy VoD services to counteract the competitive threats of Digital Broadcast Satellite (DBS). Most regard the near-term rollout of these services critical to both reducing digital churn and to increasing subscriber revenues. To date nearly all of the North American cable operators have deployed video on demand. The fantastic success of trials and early deployments has motivated cable operators to accelerate rollouts.

These service types rely on the ability of the cable networks to deliver unique, orthogonal video streams established by dynamic user control. This facility behaves in much the same way as the public telephony network; a user connects to the network (lifts the receiver), signals a unique switched communication path (dials the number) and conducts a unique session (talks to the desired party). One-to-one connectivity is necessary in telephony, and in data networks, because nearly every transaction context is unique; i.e. different content, at a different time, for a different reason. Based on conservative estimates, 10 million unique and simultaneous "one-to-one" streams will be deployed by North American MSO's in the next five years.

The move to session based video delivery will require a cost effective, high capacity switched packet network infrastructure, two way digital HFC plant and sophisticated content processing, storage and management facilities. Fortunately, the 50 billion dollars spent for HFC plant upgrades and bidirectional digital television capability underpin this endeavor. MSO's are left to focus additional spending to enable the head-end to hub network infrastructure for VoD.

Today, the cost of deploying end-to-end networks capable of delivering such a large number of independent streams continues to be prohibitive. The cost to provision a single stream (server, switch, transport, QAM, RF) is around five hundred dollars (assuming a Gigabit Ethernet based delivery infrastructure). High equipment demand, and cost pressure from MSO's coupled with technical advances and a growing competitive landscape is rapidly decreasing the per provisioned stream price.

Accelerated deployment of advanced services relies on further advancement and cost reduction of network components. This paper assumes equipment will inevitably reach acceptable price points. However, low component costs are not sufficient to enable large scale and effective VoD networks. A comprehensive architecture must be adopted to effectively provision and manage networks of sufficient scale and density necessary to support thousands of interactive and unique content sessions, a volume of sessions requiring hundreds of Gigabits per second of bandwidth.

The embodiment of such architecture must yield a system that is subjectively easy to use. It should also be scaleable in both size and capability. An optimal solution must have low capital and operational expense, high asset utilization, and manageable complexity.

EVOLUTION OF VOD NETWORKS

Early Video on demand networks deployed video servers at the edge of the network in a distributed model. The network edge is the location in the hub where the QAMs interface to the HFC plant. At the time, low utilization, scarce content and limited investment in equipment made it acceptable if not preferable to utilize resources in this fashion. In fact, some MSO's still operate large distributed VoD networks. Early deployments helped prove the business cases for VoD. In addition, important lessons about market behavior and consumer preference were learned. Experience gained during deployment of early-distributed VoD systems drove technical initiatives to further optimize VoD network design.

Distributed VoD systems are technically simple in that the delivery network is inherently localized. VoD servers sit in hubs, connected to QAM modulators and feed channel groups dedicated to serve content. Generally, servers deployed for this purpose had ASI output connections; in some cases direct QAM outputs. Asset distribution to servers varied from an intensely manual process, a technician driving from hub to hub with a TK-50 tape in his hand, to more automated systems having servers connected by an out of band channel (ATM or IP network). This channel allowed content to be inserted at a central location and copied to the remainder of the VoD servers via FTP.

MSO's rapidly realized the shortcomings of distributed VoD systems as they began to deploy en masse. In a distributed architecture there is invariably a mismatch between the playout capacity of the VoD server and the number of provisioned streams. This mismatch is due to the lack of granularity inherent in most VoD servers. The unused playout capacity drove the cost per stream unacceptably high. Storage utilization was also poor because each server in the network needed to be loaded with exactly the same content. In addition operational expenses were unacceptably high because service personnel needed to travel to each hub to maintain and upgrade VoD servers. An apparent solution was to centralize video server assets in a common head end and transport the streamed sessions to the hubs. The change to this "centralized" approach better matched server playout capacity to stream demand, thereby reducing the amount of unused server capacity. Load sharing allowed storage to be arranged such that the amount allocated to a title is proportional to the number of simultaneous sessions demanded of that title. This resulted in improved cost per stream because fewer servers were needed to service the same number provisioned streams.

ASI optical transport is routinely deployed for VoD and SVoD by MSO's today and has been an effective choice for building low to medium scale VoD systems. Rings and point-to-point ASI over DWDM networks connect VoD servers to QAM modulators in remote hubs.

There are two inherent disadvantages of ASI transport systems. The first disadvantage is that payload capacity of ASI is only about 216 MbpS. That comprises about 40 VoD streams per ASI. When deploying small-scale VoD networks with low peak load, provisioning services at this granularity is acceptable. As peak load demand increases, it is necessary to provision unacceptably large numbers of ASI links. The low ASI transport capacity undersubscribes server resources and physical transport assets (fibers, lasers etc). Replacing ASI transport with Gigabit Ethernet transport can increase the payload capacity by 500% to 1000% for the same cost. The second disadvantage is that the payload containers carried by ASI are MPEG-2 transport streams. MPEG-2 transport streams are not intended, nor capable, of forming the basis for a switched network interconnect. The virtues of a packet switched interconnect for VoD will be discussed in the next section. The effect of these two properties is unacceptable cost per bit transported. In order to reach the cost points necessary to proliferate VoD and move toward EOD, most believe that a switched packet interconnect based on Gigabit Ethernet is necessary.

GIGABIT ETHERNET AS A BASIS FOR VOD NETWORKS

The benefits of Gigabit Ethernet technology in transport systems for current and next generation VoD networks begin with its ubiquity. Quite simply, lots and lots of Ethernet equipment is bought and sold each year. This assures the contributing electronic components will remain commodity. Ethernet’s plug-and-play interoperability makes networks deployed with Ethernet easy and cost-effective to install, manage and use. Gigabit rate optical and electrical interfaces prevalent in the data communications world provide a high-speed, robust interconnect between servers, edge devices and processing elements connecting to the transport infrastructure. Contemporary servers and storage systems are optimized to operate in increments of 1 GbpS and make excellent use of off the shelf network adapters. A 1GbpS Ethernet link can transport up to 240 VoD streams, a suitable increment of streams to deploy in large scale VoD networks. Another advantage of Ethernet is its media access control (MAC) layer. The Ethernet MAC layer is a data link protocol with sufficient properties to form a sophisticated, and extremely scaleable packet switched network. Switching allows effective bandwidth management thereby reducing per stream costs.

INTRODUCTION TO NEXT GENERATION VOD TOPOLOGIES

What does the ideal VoD network look like and why? VoD networks provide service to a geographical region approximately the size of a large city. They are high in transmission capacity, moderate in complexity and provide robust, reliable interconnect. They are typically deployed as overlay networks, and are not intended to replace existing broadcast infrastructure or data/voice networks. The motivation is to develop low cost per stream networks tailored for the unique properties of VoD traffic. The systems envisioned have basis in Ethernet as the data link protocol, no different from a LAN. The departure between a standard Ethernet LAN and VoD network is the allowed distance between end terminals, intermediate path capacity, and optical route usage (DWDM). In order to develop a system optimized for VoD delivery the following capabilities must be employed.

Layer 2 Aggregation for Bandwidth Recovery

Layer 2 (Ethernet MAC layer) switching can be used to combine multiple partially filled Ethernet links to build fully utilized optical transmission paths between headend assets and hubs. Fully utilized optical links ensure each transport laser is operated at maximum capacity. System cost is reduced by virtue of needing fewer lasers. This is important because laser cost is the single largest contributing factor to overall VoD transport costs. Layer 2 capability enables aggregated links to be demultiplexed and delivered to the correct destination. Packet switching provides full mesh connection between content sources and HFC destinations allowing servers to load balance.

Figure 2 illustrates the benefits of Layer 2 aggregation over Layer 1 transport. In the Layer 1 example, content from each VoD servers can only be directed to the corresponding downstream QAM.

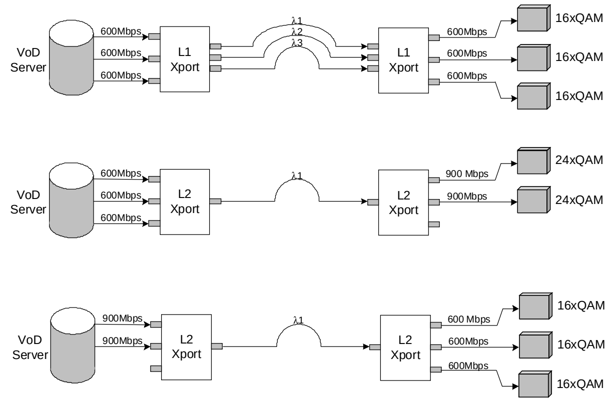

Figure 2 – Layer1 vs. Layer2 Aggregation

In the Layer 2 examples, content from all three VoD server ports can be aggregated to form a fully utilized wavelength, thus reducing fiber and laser costs. In addition video streams can be directed to any of the QAM devices from any server port on a packet-by-packet basis.

Layer 2 Forwarding and Shared Wavelength Topologies

Layer 2 forwarding allows the contruction of shared wavelength topologies. Video streams entering multiple inputs to a head end transport device can be aggregated and tagged with information which allows them to be discerned by hub end devices connected a shared optical path. Each device on the shared path can selectively receive any of the video streams available on that path and forward them to the QAM. This optimizes bandwidth utilization by allowing streams to be delivered to a number of QAMs using only the bandwidth they instantaneously require. This results in the lowest cost per video stream.

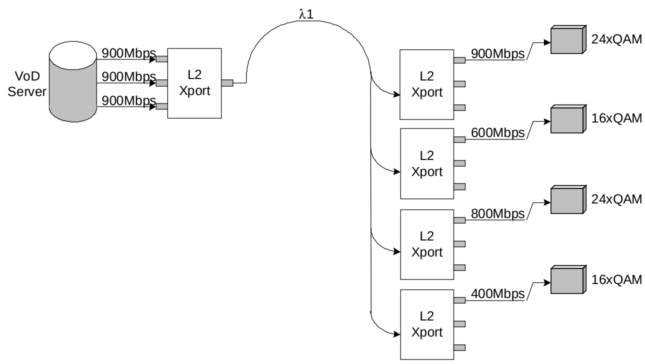

Figure 3 - Layer 2 Shared Ring Topology

An example of these benefits is evident in Figure 3. In this case Layer 2 switching aggregates three server outputs and places the video streams on a single wavelength of the shared ring. Multiple hub end devices are connected to this single wavelength. The hub end devices are configured to receive only the video streams that are destined to their associated QAMs. Ethernet allows this to be done automatically with no user intervention. Bandwidth to the QAM devices is allocated in arbitrary proportions as required. This is extremely useful in multicast environments where single copies of a video stream are presented to all hub devices, but selectively switched to QAMs based on user demand.

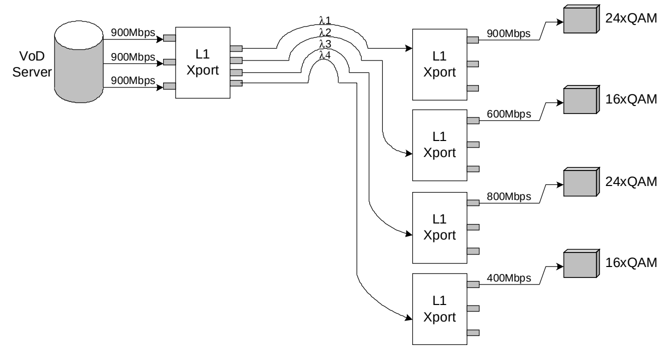

Figure 4 - Layer 1 Point-to-Point Topology

Figure 4 shows a logically equivalent topology using Layer 1. Note that the Layer 1 topology is point-to-point and requires four underutilized lasers while the Layer 2 topology with the shared ring uses only one fully utilized laser. This results in a much higher cost-per-stream compared to the Layer 2 solution.

Asymmetric Reverse Path

Up to now the topology diagrams have shown the forward video path only. There may be a need for a reverse path to carry control and management traffic from the hubs to the head end. This reverse path typically requires an order of magnitude less bandwidth than the forward path. In a Layer 1 system, bidirectional links must be used to support a reverse path. In this case the reverse path would have the same cost as the forward path even though it does not carry video content and is underutilized.

Figure 5 shows an example of asymmetric reverse path for hub to head-end interconnect. Since VoD traffic is predominantly from head-end to hub-end with comparatively low-bandwidth control traffic required for the reverse direction, Figure 5 shows a topology where a single wavelength transports the reverse path management information for all of the downstream QAMs to the head end over a single wavelength.

This preserves all but one wavelength for forward path traffic.

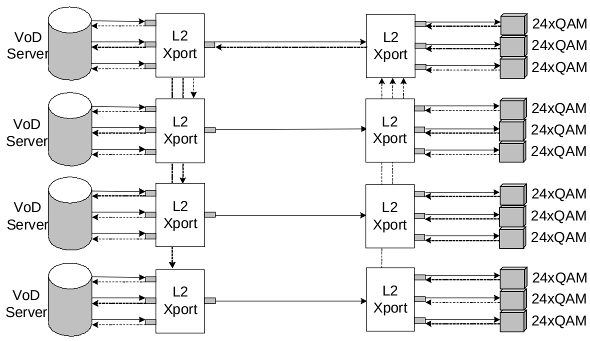

Figure 5 - Asymmetric Reverse Path

Layer 2 aggregation concentrates the reverse path control traffic from the hub onto a single gateway element, which then places this reverse path traffic onto an available optical wavelength in the ring. All other wavelengths in the hub are receive-only and do not require transmitter optics. Likewise, at the head-end the wavelengths not used for reverse traffic are transmit only and do not need receivers. This tailoring of transport optics to better match the forward and reverse path bandwidths of VoD traffic can provide significant cost-per stream savings.

The optimal VoD network makes extremely efficient use of high capacity DWDM optical components and in particular optimizes the use of cost-effective solutions available today. Efficiency is gained through fully utilizing each optical wavelength by a combination of link overhead minimization and Layer 2 traffic aggregation. Essentially each wavelength is operated near theoretical maximum capacity.

Scalability

MSO network planners require networking equipment to be cost effective and scaleable for both large and small market regions. They also require that the equipment can satisfactorily support the deployment of diverse implementations within a single large network. Furthermore, deployments may begin with a few Gigabit Ethernet links and grow to many more as the number of provisioned subscribers increases. Equipment used for the early deployments must be cost-effective for the initial small number of links yet scale and remain cost effective in the growing network.

Switched broadcast

Switched broadcast is an application intended to solve the problem of delivering limitless content choices within the finite limitations of the provisioned infrastructure by minimizing bandwidth utilization in the transport and HFC networks. By this means only program content being watched is delivered over the network. Furthermore, content watched by multiple users is transported as a single copy over the optical network. Distributed Layer 2 switching performed by transport devices in the head-end and hubs effectively utilizes bandwidth by directing content needed by that hub over a shared optical ring. Without a switched broadcast solution, video content demand will eventually exceed the available bandwidth of the cable infrastructure. Deploying systems that can support switched broadcast services which future-proof the VoD network for this eventuality.

TRANSPORT OPTICS

The optical components and topologies of the VoD transport network are designed to provide the lowest cost while fulfilling the requirements for bandwidth, reach/distance, resiliency, and ease-of-use. Figure 7 shows a typical transport network for hypothetical 5-hub model.

Video streams from the VoD servers are passed through an Layer 2 Ethernet switch, video gateway elements which perform Layer

2 aggregation of video streams onto a DWDM optical network, optical multiplexors; single-mode optical fiber to connect to an adjacent hub, optical splitters, optical protection switches, optical demultiplexors, video gateway elements performing Layer 2 deaggregation, and Gigabit Ethernet QAM devices.

Bandwidth

The VoD transport network is designed to carry VoD streams of personalized video content for each viewer from VoD servers to destination QAM devices. Each VoD stream requires 3.7Mbps of payload bandwidth quantized as 188-byte MPEG-2 video packets. Digital packet switching using Ethernet has become the defacto standard transport mechanism for VoD because of its ubiquity and low cost. The typical packet encapsulation scheme adds 3.4% overhead and uses 7 MPEG-2 packets (1316 bytes) encapsulated over Ethernet (18 Bytes) over IP (20 bytes) over UDP (8 bytes). Table 1 shows the number of VoD streams, which can be carried over a single fiber using various transport mechanisms.

Table 1 illustrates the number of Gigabit Ethernet ports provisioned based on service offering and provisioned peak load. Note that for all services above 100 Titles + SVOD and a peak load greater than 1.5%, Gigabit Ethernet links serve as ideal containers. Furthermore, 3G optics are sufficient for up to 5000 Titles while 10G optics remain highly underutilized and therefore result in higher cost per stream.

	SVOD	100 Titles	100 Titles + SVOD	1000 Titles	5000 Titles	NPVR
Peak Load% of HP	0.5%	1%	1.5%	2%	4%	10%
Streams	100	200	300	400	800	2000
#GigE's@16QAM	0.6	1.3	1.9	2.5	5	12.5
#GigE's@16QAM	0.4	0.8	1.2	1.6	3.3	8.3

Table 1 - Optical Technology Comparison for "5-Hub" Model Assuming 10% Subscription Rate

Reach / Distance

The fiber distances between the head-end node and hub nodes in the VoD transport network introduce optical loss and dispersion which need to be accounted for to maintain error-free transmission. The optical losses are introduced at connector boundaries, splices, mux/demux elements, and predominantly in the fiber itself. Network engineers compute an optical link budget by subtracting the optical receiver sensitivity and optical losses from the optical transmitter minimum output power. A link budget with adequate margin ensures adequate optical power at the receiver; while a link budget with no or negative margin indicates the need for an optical amplification device (an EDFA) in the optical path. EDFAs (Erbium Doped Fiber Amplifiers) use active Erbium-doped fiber and a laser pump source to boost the optical signal in a fiber. EDFAs become essential elements in a VoD transport network with medium to long spans between hubs.

Dispersion in optical fiber will eventually result in unacceptably low bit error rates. Because dispersion cannot be compensated by EDFAs, the optical distance (also called reach) limit is determined by the dispersion characteristics of the optical transmitter. When the VoD distance requirements exceed the laser optical dispersion limit, network architects must either add O/E/O regenerators or add another head-end, both costly alternatives.

Resiliency

Resiliency is the networks' ability to recover and sustain traffic in the presence of fiber cuts or equipment failures. VoD optical transport networks should be designed with resiliency in mind to eliminate single points of failure and provide for redundancy of critical components and a means for automatic detection and failover. The asymmetric nature of VoD and cost per stream pressures favor optical protection architectures which allow differing levels of resiliency which can be chosen to optimally balance the cost per video stream against optical protection coverage and switchover times.

The telecom industry has provided very mature network topologies (SONET/SDH and 2-fiber and 4-fiber BLSR) and elements for achieving full redundancy of optical equipment and fiber and very fast protection detection and switching times (less than 50 milliseconds) that exceed the times needed for VoD. While these can be used to provide effective resiliency for VoD transport they do so at a high cost per video stream, add additional management complexity in managing an additional transport layer, and dictate a costly symmetrical ring network topology that does not match the inherent asymmetric nature of VoD transport traffic. Unlike telecom networks, cable networks do not have strictly defined redundancy requirements and the ability to choose and optimize redundancy cost/performance is desirable. VoD transport networks that are designed for the same telecom resiliency goals (for example hybrid networks designed to carry both voice and VoD traffic), these topologies can be very effective. However, while the burgeoning VoD network infrastructure demands favor the lowest cost approach, other resiliency schemes become more favorable.

An effective strategy for optical resiliency is to design protection for those elements whose failure would affect the greatest number of VoD users. Optical protection costs for VoD can be lowered by moving the protection for fiber cuts from the Layer 1 electrical layer (as is done with SONET) to the optical layer. This eliminates the added cost of 1:1 or 1:N electrical interface protection. Ideally, individual optical transmitter and receiver redundancy protection could be optionally provided, allowing network designers to make the cost/resiliency tradeoff per wavelength while still maintaining 100% protection for fiber cuts.

Ease of Use

The operational expense of the VoD transport network is minimized and system uptime maximized by providing an integrated network which is easy to install, replicate, and maintain. The interoperability and ubiquity of Gigabit Ethernet has made the interface to the VoD server, video gateway elements, and QAM devices inexpensive and easy to maintain. The optical mux/demux, splitters, and ADPs are passive optics requiring no on-line management and have very high reliability. The EDFAs and optical switches can either be deployed as unmanaged devices which power-up in a default state or can be on-line managed; for example via SNMP or a Network Management System.

Optical Elements

The optical elements used in the transport network are well understood and have been used in existing cable networks for years. The laser transmitter performs the electrical to optical conversion of data streams. To achieve the bandwidth required for VoD, several optical channels are used per fiber using Dense Wavelength Division

Multiplexing (DWDM). The key parameters for laser choice are cost, optical reach, and overall bandwidth. 1G lasers achieve the lowest cost per laser, but the limited bandwidth results in a relatively high cost per stream. 10G lasers have high bandwidth, but are costly and have distance limitations due to optical dispersion. 10G optics also have coarse cost granularity, requiring high incremental cost as additional bandwidth is added to the network. An optimal compromise of cost per stream and optical reach can be achieved using 3G lasers with extended optical reach.

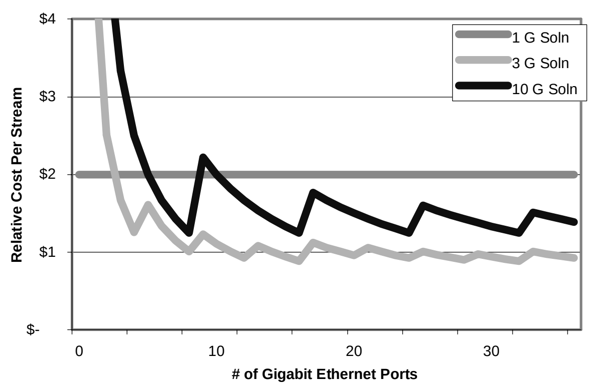

Figure 6 - Optical Laser Cost Comparison

Optical receivers are used for optical to electrical conversion of data streams. EDFAs (Erbium Doped Fiber Amplifiers) are used to provide amplification of all wavelengths on an optical fiber to offset fiber and connector losses and ensure that sufficient optical power is available for downstream receivers. Low cost passive optical mux/demux and splitter elements are used in the optical fiber distribution plant from head end to hub. Optical switches are used as protection devices to automatically switch to alternate fibers in a ring topology.

PACKET FORWARDING AND SWITCHING SCHEMES

Previous sections discussed and highlighted the subjective usefulness of Layer 1 and Layer 2 constructs and their relative benefits in developing a suitable VoD infrastructure. The following sections provide a detailed description of how traditional packet based constructs are used to fulfill the key capabilities described earlier.

Layer 1 Transport

For VoD network equipment not capable of Layer 2 and higher packet switching, Layer 1 provides a "direct wire" interconnect between the head-end and hub elements.

Layer 1 transport provides physical path connectivity between two or more endpoints on an optical link. Link information entering the input ports of a Gigabit Ethernet Layer 1 transport device are interleaved and encoded on an optical carrier. Optical wavelengths may accommodate one or more Layer 1 signals. Time Division Multiplexing generally accomplishes this. Layer 1 switching, commonly referred to as cross bar switching, offers some path flexibility by which source and destination Layer 1 ports can be cross connected arbitrarily. Layer 1 devices, however simple to provision, do not support fractional link aggregation and therefore typically underutilize optical transport capacity. For example, a Layer 1 device connected to a 600Mbps VoD server will have 60% optics and fiber utilization, while a Layer 2 device can achieve near 100% utilization by aggregating multiple fractional links. In addition, Layer 1 function does not provide shared wavelength-forwarding capabilities necessary for such schemes as switched broadcast. Layer 1 devices can only provide full duplex or simplex links. Operational complexity grows unreasonably with scale because transport assets cannot be dynamically reallocated or shared. Also, because there is no fractional rate support, asymmetric traffic patterns cannot be utilized for reverse trunking.

While Layer 1 interconnect simplifies the network topology for small-scale VoD networks, it fails to provide the benefits of bandwidth recovery, shared rings, and switched broadcast.

Layer 2 Capabilities

The Layer 2 behavior of the proposed high capacity transport network differs slightly from a conventional Ethernet LAN but is implemented in such a way that equipment costs are low and Ethernet interoperability is preserved. The following sections are a technical tutorial on the salient Ethernet functions used to achieve the desired VoD network features such as switched broadcast and shared rings.

Layer 2 classification and forwarding

Layer 2 classification and forwarding are the principal operations by which packets entering the transport domain are identified, organized and delivered to one or more destinations. The destination address contained in Ethernet packets are identified by the classification process. The result of this classification process is a list of forwarding destinations. The Layer 2 learning process ascertains the forwarding destinations. The Ethernet frame under consideration is encapsulated on the optical link and identified as a frame addressed to the downstream device listening on the selected optical wavelength. The act of identifying Ethernet frames based on their destination address, aggregating them with equivalent flows, and presenting them to downstream devices is called forwarding. By virtue of selective forwarding to interfaces based on destination information resident in the Ethernet packets, switching is accomplished. Switching forwarding and classification are used in the presented topologies to develop several classes of flows. The flow types are the following.

Path routes

Flows are grouped based on port affiliation and act like a virtual wire. Packets entering a physical port on a Gigabit Ethernet transport device are grouped together and transferred to a corresponding destination port over the fiber optic plant.

Layer 2 groups

Layer 2 groups are formed by the aggregation and dissemination of traffic to and from multiple sources. Filtering is performed and decisions are made about which packets go where on a packet-by-packet basis. The use of multicast addressing provides capability to forward the single packet to one or more destinations.

Tunneling

Tunneling is a mechanism used to trunk equivalent flows to a common end destination. Tunneling allows Layer 2 devices to forward aggregated flows on a shared virtual medium. Packets are grouped together and transported through the network as "equivalent" flows. The terminal transport node disseminates the flows through classification and delivers each packet to the described destination.

Layer 2 filtering

Layer 2 filtering works in conjunction with the classification process and provides a mechanism to scope and restrict traffic flows. Filtering can be configured to drop frames based on destination address or matched filtering criteria. A common use is to manage unknowns in the network. Unknowns are packets for which the destination is not present, or is unreachable in the network. Unknown filtering is useful in defeating broadcast storms and forwarding loops that can result in service disruption due to excessive bandwidth consumption.

Layer 2 add, drop, pass

Layer 2 add, drop, pass capability allows the formation of packet rings and provides a basis for multicast and broadcast over a single wavelength. Layer 2 flows are injected and terminated by members of the optical ring. Multiple members of the ring can receive a singular flow. Shared optical wavelength paths, used in conjunction with "star over ring" paths, provide a way to organize transport bandwidth based on steady state and transient load. Star topologies deliver the basis bandwidth which is the steady state load. Ring topologies provide a mechanism for bandwidth leveling and a shared medium for delivery of content assets, and switched broadcast services. This hybrid approach offers the ability to manage and mitigate transient bandwidth demand with little or no over provisioning. Both star and ring connections may exist within the same fiber or within tunnels on the same wavelength

Layer 2 Path and Flow Aggregation

Layer 2 Path and Flow Aggregation are used to concatenate traffic and fully utilize optical paths. They also provide the ability to utilize aggregation to virtual domains for transport of more granular traffic flows such as reverse path traffic. Asymmetric reverse path forwarding is an example of Layer 2 flow aggregation that exploits the traffic patterns in VoD transport networks.

Layer 2 asset provisioning and load balancing

A significant advantage gained through Layer 2 switching is the ability to ideally match content delivery assets to transport infrastructure capability. Switching allows equipment providing session fulfillment, such as video servers, to reach any endpoint in the network. This allows servers to share the workload in delivering VoD sessions.

Layer 2 address learning and aging

This function utilizes Ethernet source address awareness and classification to determine destinations for packets impinging the network (i.e. you don't have to know which downstream port to plug the QAM modulator into learning will find which port its plugged into). Learning also provides information to forwarding/filtering functions to automatically determine which end points are reachable by the network.

SUMMARY

The VoD architectures presented in this paper can be summarized by considering the high capacity 5-hub 300,000 homes passed VoD network model presented earlier. The key architectural and cost advantages of this model are high VoD stream capacity per fiber, long optical reach, automatic optical protection switching, and a low-cost asymmetric reverse path.

The optical transport system is constructed as a hybrid ring/star architecture. A northbound ring and a redundant southbound ring connect the head-end and five hubs. Each ring has 100GHz spaced

DWDM 3Gbps wavelengths, which in aggregate carry, over 33,000 VoD streams. Transmitter/receiver optical link budgets exceeding 32dB are achievable with low-cost laser transmitters and six spans each exceeding 25km without the need for regeneration and using two EDFAs per ring. A lower cost per stream point could be achieved with shorter reach optics, but would add more network design complexity and limit the flexibility of deploying the same architecture in nearly all VoD deployment areas.

The duplicate rings provide 1:1 protection for fiber cuts anywhere along the transport ring as well as failure of any of the optical splitters and EDFAs used on the ring. By monitoring the recovered optical power at the destination, the optical protection switches can failover without any user intervention. The topology can be constructed with an asymmetric reverse path where the head-end uses mostly transmitter-only optics and the hub-end uses mostly receiver-only optics. This reduces the optical transceiver components (which can be 20% to 40% of the overall transport network cost) nearly in half.

In the hybrid ring/star topology, the Ethernet switch at the head-end forms a star Ethernet connection between the VoD servers and the QAMs. The gateway elements then aggregate the VoD streams onto an optical wavelength, each of which is connected via the optical transport ring. A corresponding QAM at a hub-end demultiplexes the VoD streams from the optical wavelength and switches individual VoD streams to the appropriate QAM device. This allows any VoD stream from any VoD server to be directed to any head-end QAM output.

Gateway elements vary in complexity from simple Layer 1 devices to fully featured Layer 1-4 devices capable of advanced features such as switched broadcast. A Layer 1 device functions as a “wire” connecting entire Gigabit Ethernet ports (with all VoD streams remaining intact) from the head-end Gigabit Ethernet switch to a hub-end QAM. Adding Layer 2 capabilities into the gateway elements allows individual Ethernet frames (and the VoD streams they are carrying) to be individually switched and aggregated which gives much finer granularity in provisioning traffic from the VoD servers to the hub-end QAMs. This allows advanced features such as traffic aggregation and load balancing which can save VoD costs by reducing the required number of Gigabit Ethernet switch and QAM ports.

In conclusion for a VoD transport network to be cost effective and scaleable, it must be more than just a large pipe. Optimal use of fiber bandwidth and reducing the cost of transport optics are essential to reducing per stream costs. Layer 2 VoD stream aggregation reclaims fallow bandwidth. Asymmetry in the reverse optical path reclaims fallow fiber bandwidth and cuts the cost of DWDM transmitters and receivers nearly in half.

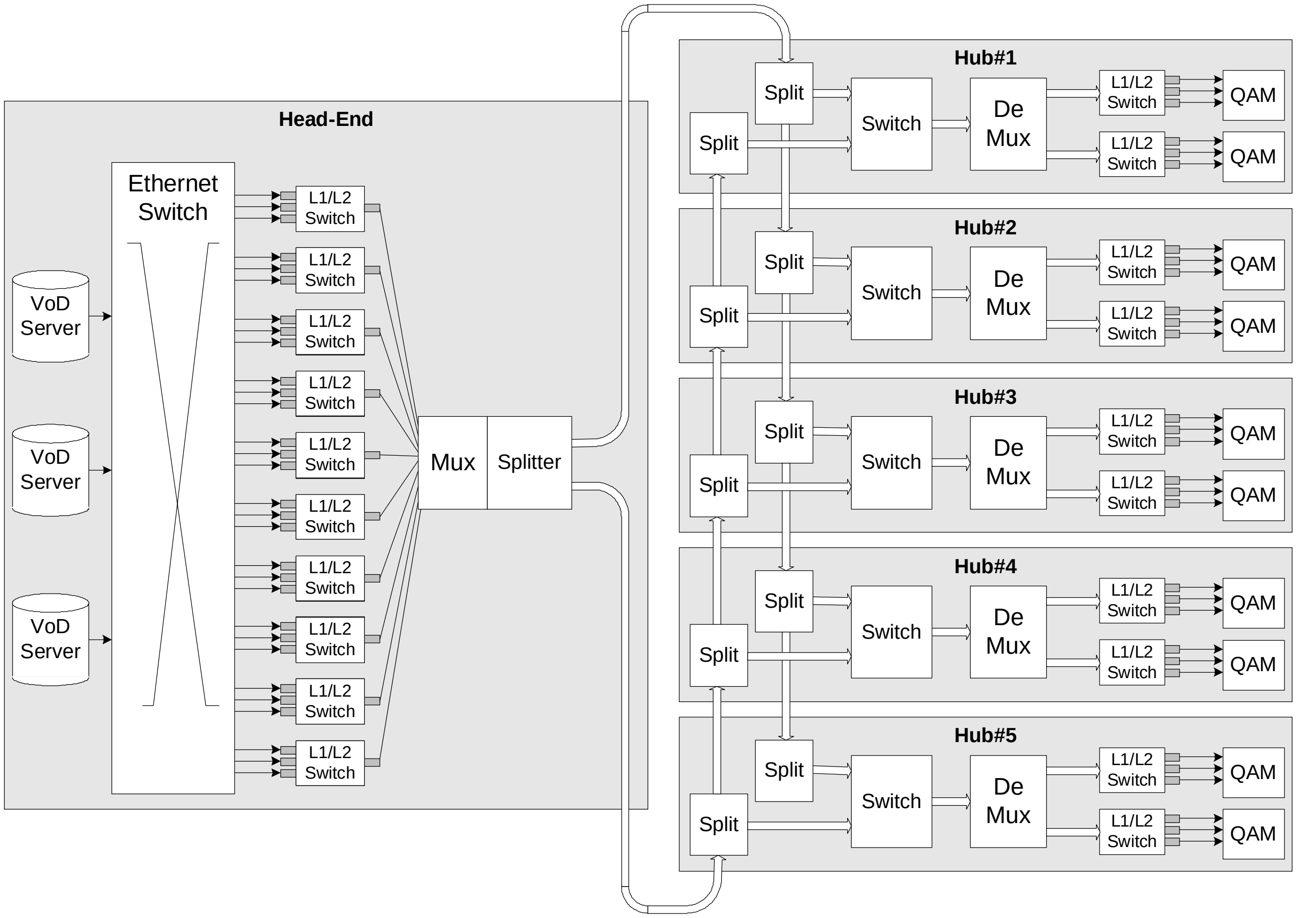

Figure 7 - 5-Hub VoD Network Model

CONTACT INFORMATION:

John Amaral
Artel Video Systems
237 Cedar Hill Street
Marlborough, MA 01752
jamaral@artel.com
508-303-8200 x285

Paul Pilotte
Artel Video Systems
237 Cedar Hill Street
Marlborough, MA 01752
ppilotte@artel.com
508-303-8200 x304