WHITE PAPER

Video Transport and Distribution for IPTV Networks

Matt Hallam — Senior New Business Development Manager

Tom Rarick — Senior Principal Engineer

Executive Summary

IP video can place a high demand on network infrastructures, especially as the number of on-demand services and high-definition subscribers grows over a large geographical area. Ensuring the best user experience requires a network infrastructure that flexibly adapts to changing subscriber demands. As bandwidth demand increases, fiber is pushed out to the network edge toward the subscribers. This allows adequate capacity for growth in consumer video applications, as well as enterprise services, wholesale services and network infrastructure consolidations. Transporting and distributing very high bandwidth services such as video is a job well-suited for the Tellabs integrated transport and flexible access infrastructures.

Introduction

Providing new IP-based entertainment requires that service providers design new networks capable of delivering broadcast (linear) and on-demand (non-linear) services to a multitude of in-home devices. These networks must be able to scale dynamically beyond the limits of today's requirements, but with the reliability equal to the five 9's networks of today. However, traditional network design and deployment methodology presents operational and financial challenges that could significantly impact the business case. Overcoming these challenges requires service providers to deploy virtual service networking that is efficient and scalable over an integrated transport infrastructure.

Given the emergence of internet-based video offerings from major television networks, including downloadable videos and on-demand service from Google, Yahoo, Apple and major studios, some service providers are concerned that their networks will be used to access content that competes with their own IP video offerings. As service providers contemplate significant investments in network upgrades to enable IP video delivery, they likely wonder whether these investments will pay off. Rather than seeing this new infrastructure as a threat, service providers should use the increased bandwidth to their advantage and deliver compelling new services to their customers. At the same time, the new infrastructure must continue to support legacy services to maintain existing customers and revenues.

The very nature of the IP video traffic will change over time. IP video has to come from an intelligent infrastructure that offers service providers flexibility, growth and reliability. Initially, traffic may be generated by broadcast or multicast services, but TV viewing behavior is shifting away from linear viewing toward time-shifted or non-linear viewing. As showin in Figure 1, traffic will be increasingly dominated by interactive services, such as Video on Demand (VoD) and Network-based Personal Video Recorders (nPVRs).

Exactly how video networks evolve in the future may be uncertain, but it is possible to plan for it. To prepare for consumer changes, service providers must have an extremely flexible infrastructure in place that will serve both the broadcast market of today and consumers' requests of the next 5 to 10 years. This infrastructure shouldn't force the service provider to commit to a single technology.

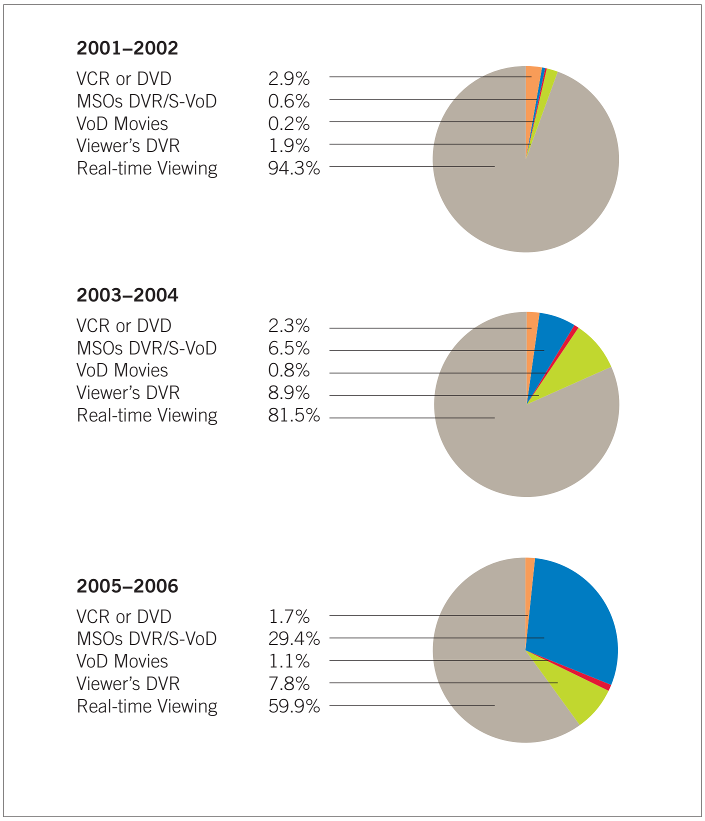

Figure 1. Evolution of subscriber viewing

The ideal infrastructure is one that has been designed from day one to enable service providers to:

- Rapidly add new service distribution elements
  - On-demand
  - Localized ad insertion
  - nPVR
- Scale the network in all dimensions
  - Bandwidth capacity
  - Number of subscribers
- Deliver all services reliably and quickly
  - Quick channel changes
  - Very low downtime

Interactive services, such as VoD, nPVR and HDTV, will place different service delivery requirements on the network infrastructure. The sheer volume of simultaneous consumer interactive sessions will likely become an issue in itself. And, while service providers' upgrade cycles can take many years, consumer technology and service demands evolve more rapidly. Who could have predicted the success of the MP3 player just five years ago? As service providers plan to upgrade downstream bandwidth to the consumer, consumers increase their demand for upstream bandwidth. It is not clear what new applications will arise in the next few years. Rather than make technology decisions that create service barriers, today's choices must enable tomorrow's service delivery opportunities.

Video Transport Challenges

Video distribution represents a significant change to the optical infrastructure. In the world of voice and internet service, the average bandwidth per subscriber is on the order of 100 Kbps. Even with Digital Subscriber Lines (DSL) rates up to 6 Mbps, the utilization is much lower — as much as a factor of 200 lower.

Video services can be broadly categorized as either broadcast or unicast. Broadcast represents the channels to which all subscribers have concurrent access, including the standard channel lineup and services like pay-per-view in which programming is transmitted at scheduled times.

Conversely, unicast services represent video content that is streamed uniquely to a subscriber, including VoD services. Interestingly, network-based PVR functionality also represents a unicast relationship with the subscriber.

With standard definition bandwidths of 3.75 Mbps and high-definition bandwidths of 10–18 Mbps, broadcast content alone can easily total 1 Gbps with 200 standard definition channels and 20 high-definition channels. If the HD content increases to 100 channels, 2 Gbps would be required for broadcast content alone. The bandwidth requirement could easily increase to 3–4 Gbps with more channels and more HD content. The bandwidth total is multicast to all households in some manner, so the effect to the video transport network is at least 4 Gbps, up to 8 Gbps in a maximum use scenario.

VoD bandwidth, however, is a different story. VoD bandwidth requirements are likely to be small at the outset as consumers maintain traditional television viewing habits and initial take rates are low. However, there are a number of drivers that could cause the demand for unicast video traffic (i.e., traffic targeted to one Set Top Box (STB) or television) to skyrocket, including:

- Consumers become more comfortable with time-shifted TV viewing through the increased use of PVRs
- Increased take rates for VoD services (whether purchased individually or as part of a package)
- Carriers deploy network PVR capabilities that allow subscribers to rewind and fast forward a particular program — this approach may actually be preferred over home-based PVRs to address piracy concerns and preserve digital rights

Because these reasons apply to all televisions in a household, the bandwidth required to accommodate unicast services will likely grow to very large numbers. For example, an end office serving 10,000 subscribers could see 6 Gbps of unicast video with only 10% of the subscribers watching VoD programming. This shift, in which unicast bandwidth requirements exceed broadcast bandwidth requirements, represents a significant move from a push model to a pull model for video delivery.

As unicast traffic increases, the hubbed architecture for broadcast services may give way to a more distributed architecture in which the video network edge moves closer to the subscriber. VoD servers and network PVR functions pushed to the edge would allow the network to scale as necessary. Thus, the challenge for the service provider is to position the transport and access infrastructure cost-efficiently today, yet plan for graceful migration to accommodate the inevitable changes that will push video-related processing closer and closer to subscribers in the future.

Video Distribution Today

As shown in Figure 2, a major portion of today's video distribution networks represent an evolution of traditional DSL networks for Internet service. The goal is to add significant bandwidth to the Digital Subscriber Line Access Multiplexers (DSLAM) while partitioning the intelligence enough to preserve the DSLAM architecture. Additionally, to conserve transport bandwidth, the complete set of broadcast channels only reach the edge router at the Video Hub Office (VHO), not the end office.

To achieve this structure, multiple layers of multicast — from the edge router to the local Ethernet switch to the DSLAMs — imply that the highly dynamic activity of channel changes are distributed throughout the network, affecting the DSLAM, the Layer 2 aggregator and the edge router.

Another large portion of current video transport networks are variations of Synchronous Optical Networking (SONET) infrastructures. In the typical regional network for video transport shown in Figure 3, one or more long-haul Points of Presence (POP) locations within a region provide access to broadcast video satellite farms at the Super Head-Ends (SHE). SHEs may be physically located significant distances from the regional networks — in fact there is likely more than one SHE across the carrier's network to provide tolerance against faults and sunspots. The broadcast video from these remote locations is delivered to the regional network and to the VHOs as a set of GbEs over virtually concatenated SONET transported over legacy Wavelength Division Multiplexing (WDM) transport.

The VHO is the main video distribution office within the regional network. It is the top of the tree for multicast networks, the source for VoD traffic and the entry point for local broadcast content and local advertisements. From the VHO, the network is actually a large tree network with three or more layers of distribution ultimately ending at the subscriber. Out from the VHO are aggregation offices, which are merely convenient locations within the carrier's network to distribute traffic to and from the end offices. The end offices contain the equipment that ultimately distributes video to the subscribers.

The transport between the POPs, VHOs and aggregation offices is typically performed by point-to-point WDM with regeneration at every office. The transport from the aggregation offices to the end offices is typically via stacked SONET OC-192 rings. On any given ring, capacity may be dropped at an end office, requiring a SONET Add/Drop Multiplexer (ADM). If capacity isn't dropped at an office, a SONET regenerator is required. As the need for capacity in the end office increases, more SONET equipment must be installed and

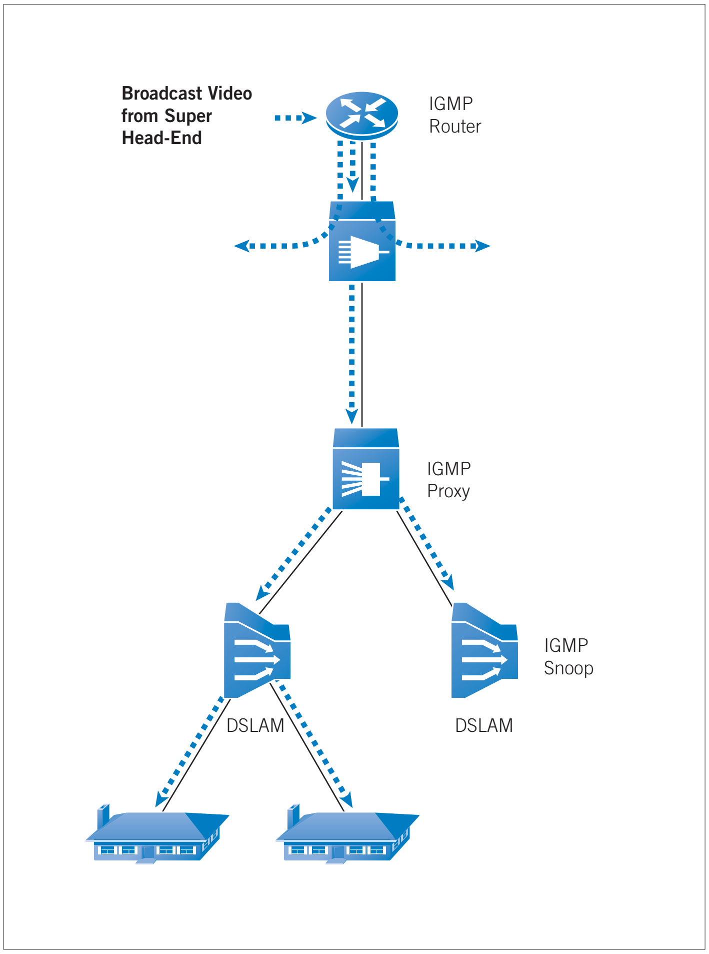

Figure 2. Typical video distribution architectures

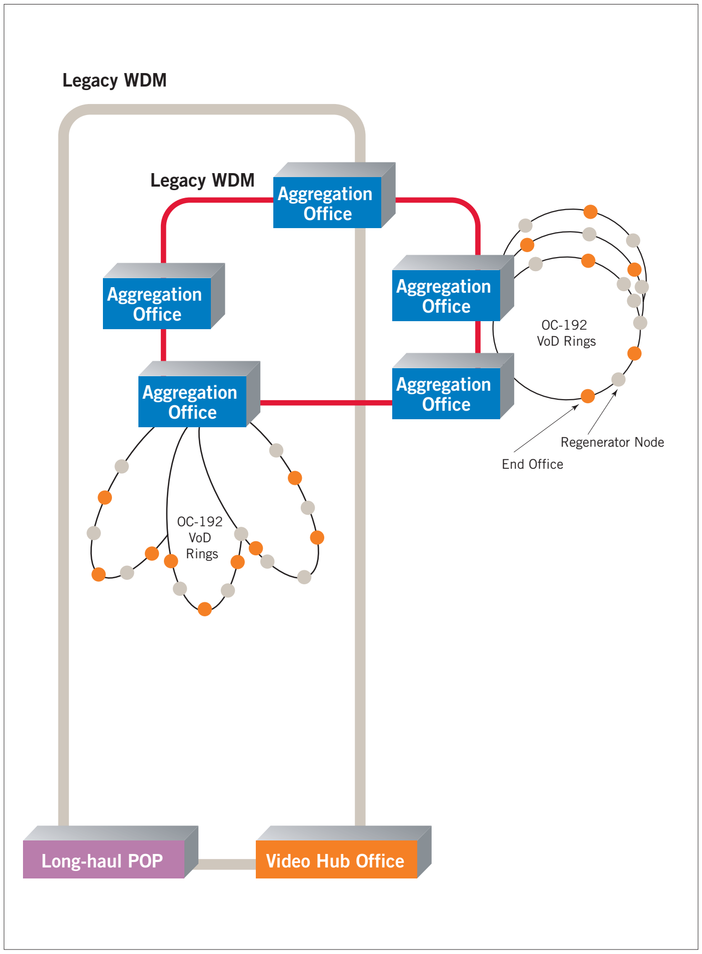

Figure 3. Typical video transport network

when the ring is exhausted, an additional ring is installed in parallel. It's easy to see how the anticipated demand for sophisticated video services will have a dramatic effect on bandwidth requirements, a corresponding increase in the number of SONET nodes and rings, and a proportionate rise in CapxEx and OpEx.

Though SONET has many great qualities, it is not the optimal solution for delivering gigabit-level packet services like IPTV. SONET can transport Gigabit Ethernet via virtual concatenation, but it is not the most cost-efficient means for transport. And, as all services move to packets, the rigid hierarchy of SONET will require expensive intermediate multiplexing points just to access the service.

These are industry-recognized issues that are being addressed by the International Telecommunication Union (ITU). Work is already well underway to develop a packet-based transport mechanism featuring all the strengths of SONET, but optimized for packet-based services. The solution — MPLS transport — will be discussed in more detail later in this document.

In summary, building video transport and distribution networks today is a very expensive proposition — multiple high-touch boxes are costly to buy and operate. Existing SONET transport solutions simply are not scalable to the size necessary to handle the demand, nor are they optimized for gigabit packet services. All of these factors drive up CapEx and OpEx dramatically. From both technology and budgetary perspectives, a new solution better aligned to the dynamic, high-bandwidth nature of video transport is required.

The Tellabs Flexible Access & Transport Solution for Changing Video Demands

The strategic evolution of video transport and distribution is a fundamental driver for the Tellabs end-to-end multiservice architecture (Figure 4). The demand for bandwidth at the edge pushes fiber closer to subscribers and extends DWDM to end offices. The result is adequate capacity for consumer video applications, enterprise and wholesale services, and network infrastructure consolidations.

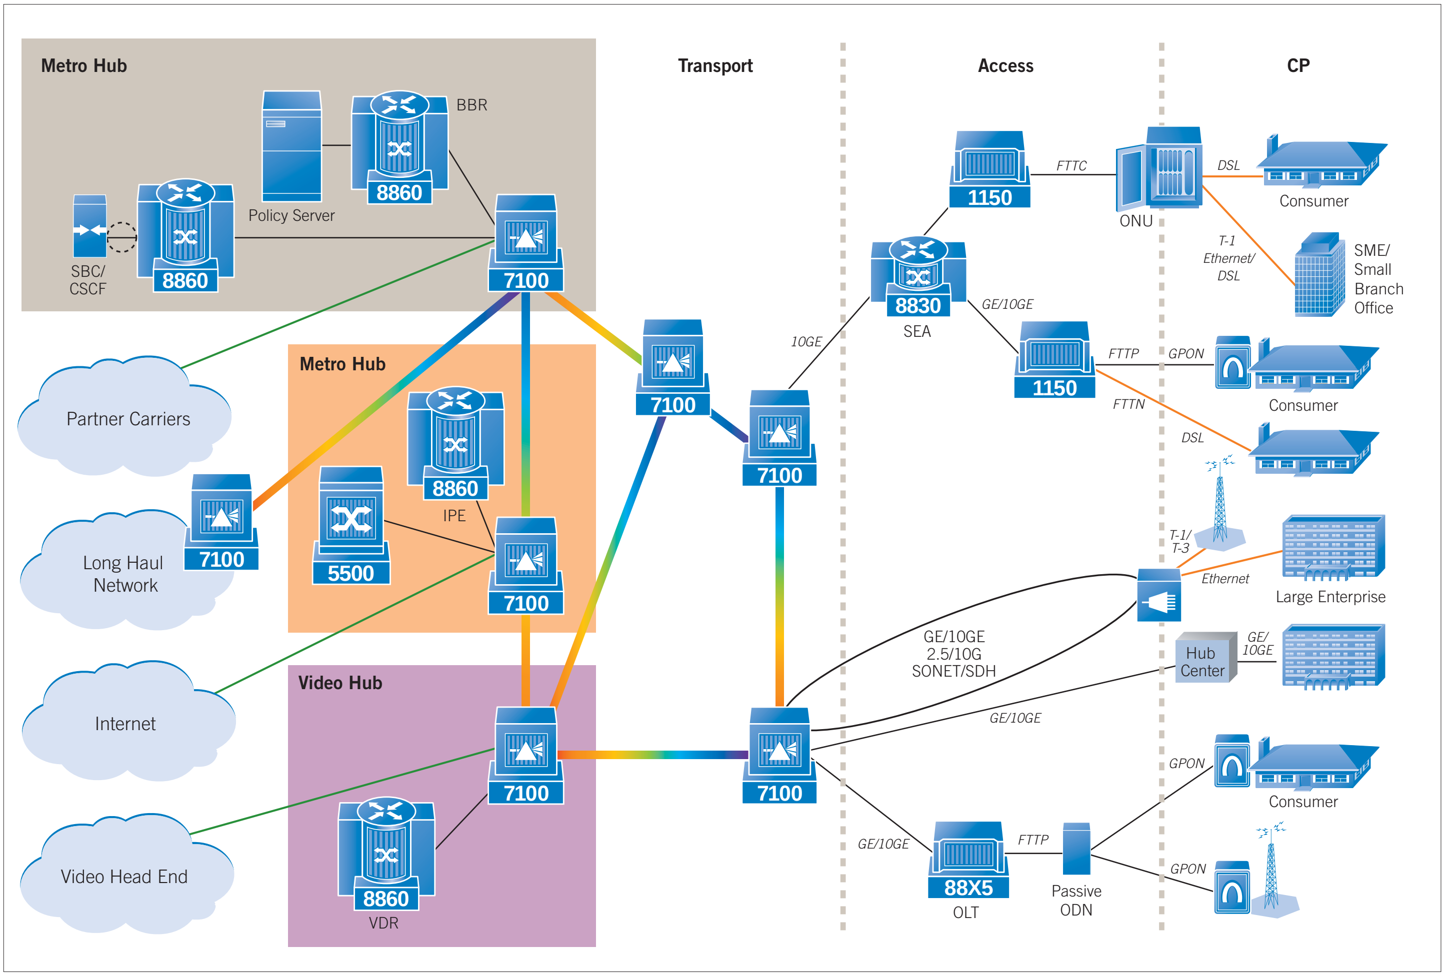

Figure 4. Tellabs end-to-end architecture

The Tellabs end-to-end architecture:

- Provides an integrated transport and access infrastructure via optical and packet transport with fiber pushed deeper into the access network
- Pushes subscriber awareness to the edge of the network
- Supports QoS for video, independent of other services that may also be running on the converged infrastructure, via traffic-engineered Logical Transport Topologies (LTTs) that enable the video service portion of the transport and access infrastructure (or any other service) to be logically "managed"
- Enables graceful evolution from legacy services to new services through support for TDM transport and other legacy services

Transporting very high bandwidth services like video is a task well-suited to Reconfigurable Optical Add/Drop Multiplexers (ROADM), a key component of the integrated transport architecture. As illustrated in Figure 5, ROADMs provide the capacity to interconnect core offices, such as the long-haul POPs, VHOs and aggregation offices. ROADMs also support flexible and dynamic switching at the optical layer, enabling quick and easy modifications to network designs as demands change. Optical pass-through capailities also enable ROADMs to be utilized as inexpensive mechanisms to bypass offices — a practical alternative to SONET ADMs and regenerators.

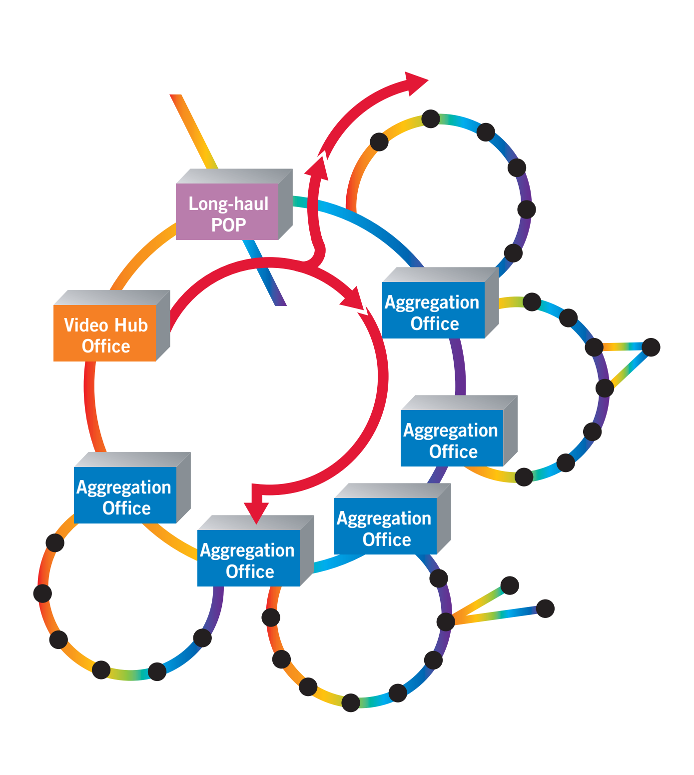

Figure 5. Video transport over a ROADM network

"You hear a lot about servers and set-tops, but ROADMs are the enabling technology that really make IPTV possible."
G. Keith Cambron
Senior Vice President, AT&T Labs
TelephonyOnline — March 20, 2006

ROADMs also deliver optical drop-and-continue capabilities, enabling the distribution of optical channels to multiple subnetworks and nodes with a minimum of lasers. For example, in a typical regional video transport network, broadcast channels can be distributed to all end offices with only one laser, providing a very low-cost way to distribute video to dozens of offices. This capability also supports very high-quality distribution as video bypasses intermediate routers, minimizing jitter and delay. Bypassing routers also reduces the costs required to acquire and maintain additional router ports.

With their integrated amplification features, ROADM-based networks can extend very long distances — up to 1,000 km — without requiring costly regeneration. This capability allows carriers to completely rethink optical network architectures as rings can contain many nodes and be much larger. Existing ROADM rings can be extended to reach larger coverage areas, including spurs.

"Overall, ROADM-based optical transport networks can provide savings of 50 percent or more over the SONET alternative."
"Metro Network Reroute" by Sr. Writer Ed Gubbins
TelephonyOnline — March 20, 2006

Video represents one of several Virtual Service Networks (VSN) supported by the Tellabs end-to-end architecture. Figure 6 offers a functional view of the video distribution network, showing more detail in the end office. The top of Figure 6 is a flattened view of the physical network; the middle portion illustrates a functional view of the broadcast network; and the bottom of the diagram features a functional view of the unicast video network.

At the VHO, Video Processors (VP) provide local content and ad insertion for broadcast traffic, and Video Servers (VS) provide VoD content for unicast traffic. These sources then connect to the Video Distribution Router (VDR) for transport over the optical transport network to the end offices. Although the SHE is shown, it doesn't play a role in unicast flows, which actually begin in the VHO.

For broadcast flows, optical drop-and-continue within the ROADM provides cost-effective traffic transport to the end offices. The ROADM also includes Ethernet multicast functionality to replicate the broadcast video streams within the end office. For unicast flows that are point-to-point in nature, the traffic is carried via packet transport trunks using Transport MPLS (T-MPLS) functionality integrated into the ROADM. These T-MPLS trunks are traffic-engineered to the planned capacity at the end office, allowing a simple Call Admission Control (CAC) function at the VHO to determine if there is capacity in the network for any given unicast session.

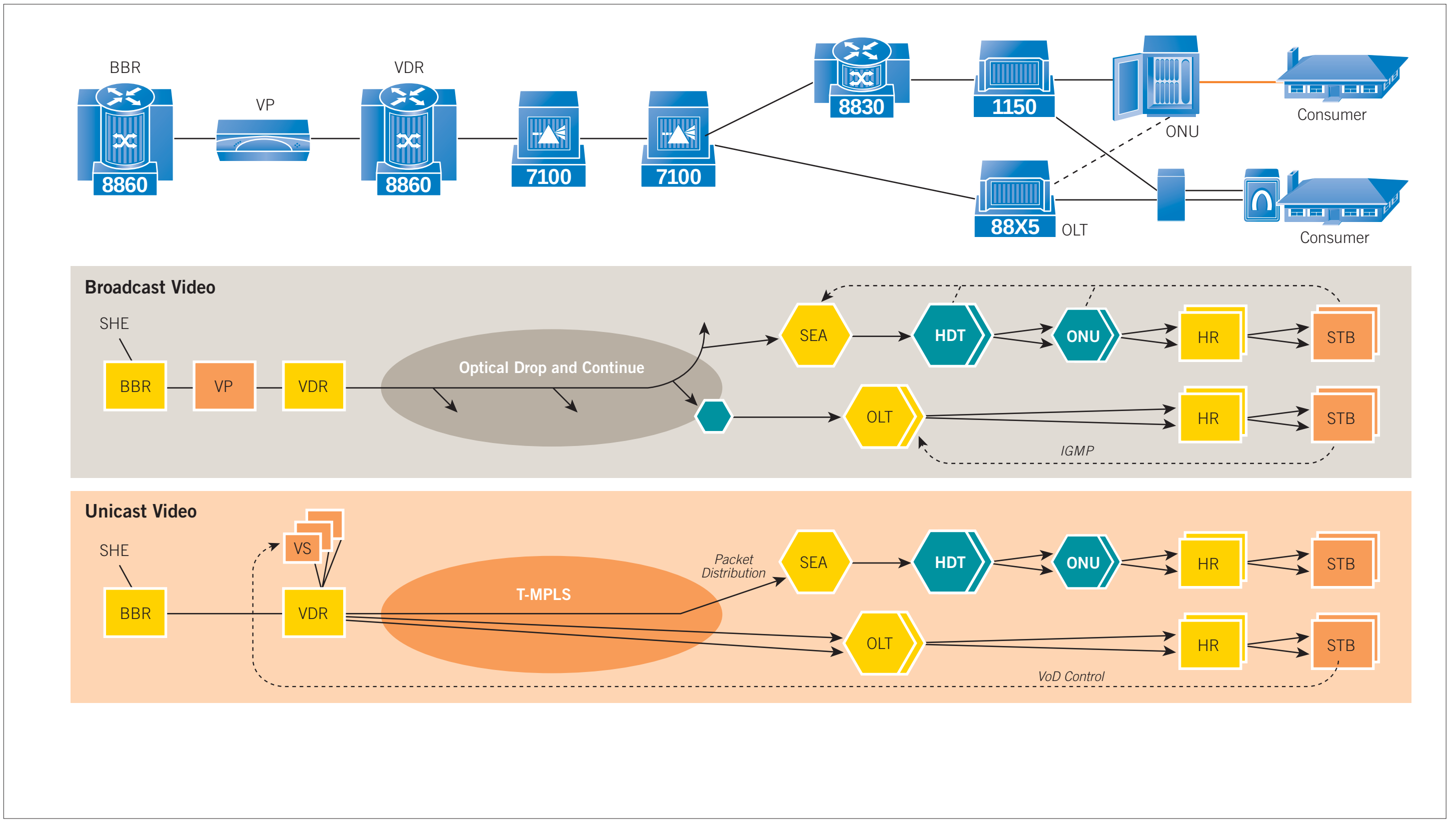

Figure 6. Tellabs video distribution architecture

At the end office, Tellabs offers two alternatives to meet carriers' specific needs. The Tellabs Multiservice Access Platform (MSAP) is a flexible and cost-effective video distribution solution featuring a comprehsive set of available access technologies:

- Fiber-To-The-Premise (FTTP) with Gigabit Passive Optical Networking (GPON)
- Fiber-To-The-Curb (FTTC) with VDSL2
- Fiber-To-The-Node (FTTN) with ADSL2+ and VDSL2

As subscriber bandwidths increase to 50 Mbps or more, the deep fiber methods of FTTP and FTTC become the preferred methods for access. The architecture ilustrated in Figure 6 features the FTTC solution on the top path, with the Optical Networking Unit (ONU) pushed very near the subscribers. The ONU provides Layer 2 aggregation for the Head-end Digital Terminal (HDT) located in the end office, and the HDT provides further Layer 2 aggregation for the Service Edge Aggregator (SEA).

The SEA enables per-subscriber functions and aggregates like services from multiple subscribers to common VSNs before transport to the service delivery point. In the case of VoD, all traffic destined for an end office from the VS uses a common VoD-based VSN.

The second distribution alternative — optimized for FTTP-only access — is a low-cost alternative to deliver high-bandwidth subscriber services over fiber. An Optical Line Terminal (OLT) delivers very high bandwidth and its associated performance to subscribers via GPON. The OLT provides a combination of the HDT and SEA functions described in the MSAP model, aggregating all subscriber traffic and assigning it to the appropriate VSN. Serving as the policy enforcement point of a policy-based network, the OLT also perfroms per-subscriber awareness functions, including policing and shaping.

The FTTP solution also provides the horsepower for new services, increased bandwidths and very fast processing. With this architecture, all broadcast channels are sent to the FTTP OLT, eliminating the complicated typical multicast architecture illustrated in Figure 2. Internet Group Multicast Protocol (IGMP) requests from subscribers are handled efficiently by the OLT alone.

The integration of the SEA into the OLT provides significant savings when compared to the deployment of separate devices, as shown in Figure 7. By integrating the T-MPLS and Ethernet functionality into the ROADM, service providers can save significantly over deploying stand-alone Ethernet switches. These savings could easily exceed 30 percent in CapEx alone, with additional OpEx savings gained due to the reduced touch points in the network.

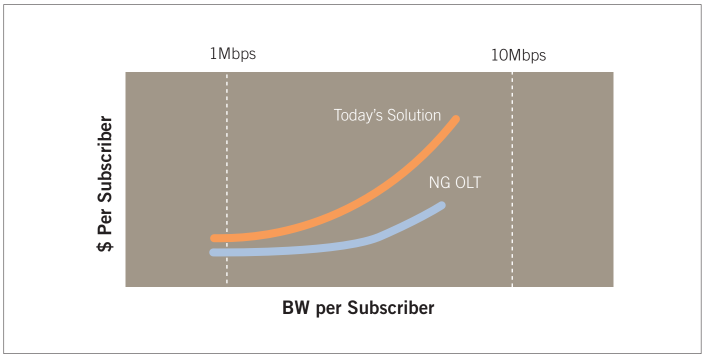

Figure 7. Integrated OLT and SEA cost savings

Conclusion

Service providers face significant challenges associated with the delivery of today's video services — challenges that will only become more daunting as broadband market needs expand in the future. SONET and DSLAM-based access solutions do not provide the scalability or desired cost points necessary to deliver cost-effective high bandwidth subscriber services. Rapid developments in technology and the demands of sophisticated consumers require a nimble video transport and delivery infrastructure that can quickly accommodate both network growth and consumer demand for new and more sophisticated services.

The Tellabs end-to-end architecture — featuring ROADM-based video transport and versatile access solutions — offers flexibility, scalability and reliable technology to deliver high-qualitiy subscriber services via a highly cost-effective network.

North America

Tellabs
One Tellabs Center
1415 West Diehl Road
Naperville, IL 60563
U.S.A.
+1 630 798 8800
Fax: +1 630 798 2000

Asia Pacific

Tellabs
3 Anson Road
#14–01 Springleaf Tower
Singapore 079909
Republic of Singapore
+65 6215 6411
Fax: +65 6215 6422

Europe, Middle East & Africa

Tellabs
Abbey Place
24–28 Easton Street
High Wycombe, Bucks
United Kingdom
HP11 1NT
+44 870 238 4700
Fax: +44 870 238 4851

Latin America & Caribbean

Tellabs
1401 N.W. 136th Avenue
Suite 202
Sunrise, FL 33323
U.S.A.
+1 954 839 2800
Fax: +1 954 839 2828

74.1716E Rev. A 10/06