USING LAYER 1 MULTICAST OVER DWDM WITH COUNTER-PROPAGATING WAVELENGTHS FROM TWO NATIONAL MASTER HEADENDS

Lead Author: Mark Pellegri, Senior Manager Transport Networks, Cox Communications

Co-Authors: Joel Bethers, Cox Communications; Mike Gricus, Cox Communications; Gaylord Hart, Infinera Corporation; Brian West, Cox Communications

EXECUTIVE SUMMARY In 2010 Cox made a fundamental change in the way national broadcast video content is distributed to its serving markets. This change was dubbed the "V2B 2.0" initiative and involved using dedicated wavelengths, constructed in a tree-and-branch manner, to deliver the content from two master national headends rather than sending the content over the same routed Layer 3 network used to carry other services such as High Speed Internet, telephony, entertainment on demand, signaling traffic, and business Ethernet services.

Since its launch Cox has experienced an array of reliability, engineering and operational benefits directly related to the new architecture. This paper will explore the evolution of Cox's video distribution models, the details of the V2B 2.0 model's architecture and some specific OEO capabilities of the DWDM platform used in our national transport backbone and how they enabled the V2B 2.0 model to be such a success. Lastly, some challenges that remain unsolved by V2B 2.0 will be illustrated along with a brief discussion on why the V2B 2.0 model that is so effective on backbone networks is not necessarily the best choice within the market metro networks.

Video Distribution History

The distribution of linear broadcast video programming to Cox's operating systems has evolved through three distinct phases: local content acquisition, centralized acquisition with distribution to the systems via Layer 3 multicasts, and centralized acquisition with distribution via Layer 1 multicast. The latter two phases are referred to internally as "Video to the Backbone" or V2B 1.0 and V2B 2.0, respectively.

Local Headend Acquisition

The acquisition of national broadcast video content in all Cox markets previous to launching V2B (Video to the Backbone) in 2007 was performed at local headends in each market. This method was the only feasible way of obtaining content due to the nature of our national data backbone at that time. In this era the Cox national network backbone was built upon leased circuits from various national carriers. The costs of these circuits was directly related to the amount of bandwidth needed, as well as distance traversed, so trying to distribute vast quantities of programming channels across the national backbone was cost prohibitive. Further, the lead time for establishing new circuits often reached six to nine months from the time of order. This lead time made it difficult to keep up with residential internet demands, let alone the rapid pace of video programming expansion, which was already approaching 2 Gbps for national broadcast content.

In general two methods were used to acquire national content at the individual market headends: via satellite dish “farms” or via direct fiber feed (typically for either local stations or when national studios happened to be located in the same market.) In some cases antennas were also used for local station acquisition. Both of these methods were acquired and transported throughout the market based on widely established RF formats. The fiber transport technologies were typically of a proprietary nature developed by video equipment vendors (e.g. Motorola, Scientific Atlanta, etc.) and didn’t necessarily utilize common lower layers such as IP, which further confounded efforts to distribute channel acquisition beyond one market. Significant local staff resources were required to maintain various satellite dishes, proprietary transport technologies, transcoders, statistical multiplexers, etc. as well as to negotiate locally with national content providers for transmission rights, procure and build the proper receivers, and perform various other duties as required to maintain the infrastructure. This model resulted in a very large variance in the equipment deployments and channel lineups across the entire Cox footprint and made it quite difficult and costly to share any local content across markets at any type of scalable level.

Until the advent of V2B each market built a transmission system that stood apart in an “unconverged” fashion from any other service or network within the market (e.g. telephony or High Speed Internet service.) The video transmission system could function quite nicely without any outside involvement even with new services such as Video on Demand (VOD) due to the content being acquired within the market. The only exceptions to this would be for vendor remote access or communication between the billing system and the DNCS or DAC. This model worked reasonably well, albeit expensively, until Cox started to see increased competition and market pressures that drove even greater channel lineups—HD in particular—that forced Cox to rethink the approach to acquiring content in terms of service velocity (scalability) and ubiquity across the Cox footprint.

V2B 1.0: Central Acquisition at Master Headends, Distribution via Layer 3 Multicast

V2B 1.0 was born out of a need to meet the demands of increased competition to offer a rich variety of HD content. This large-scale project was initially targeted at being able to provide added HD content to every market in the Cox footprint while controlling costs and resource scalability. The content would be acquired at two geo-redundant locations, encoded, stat-muxed, and then delivered over the Cox national Ethernet IP backbone to each market. A key enabler of V2B 1.0 was the onset of Cox’s own national DWDM transport

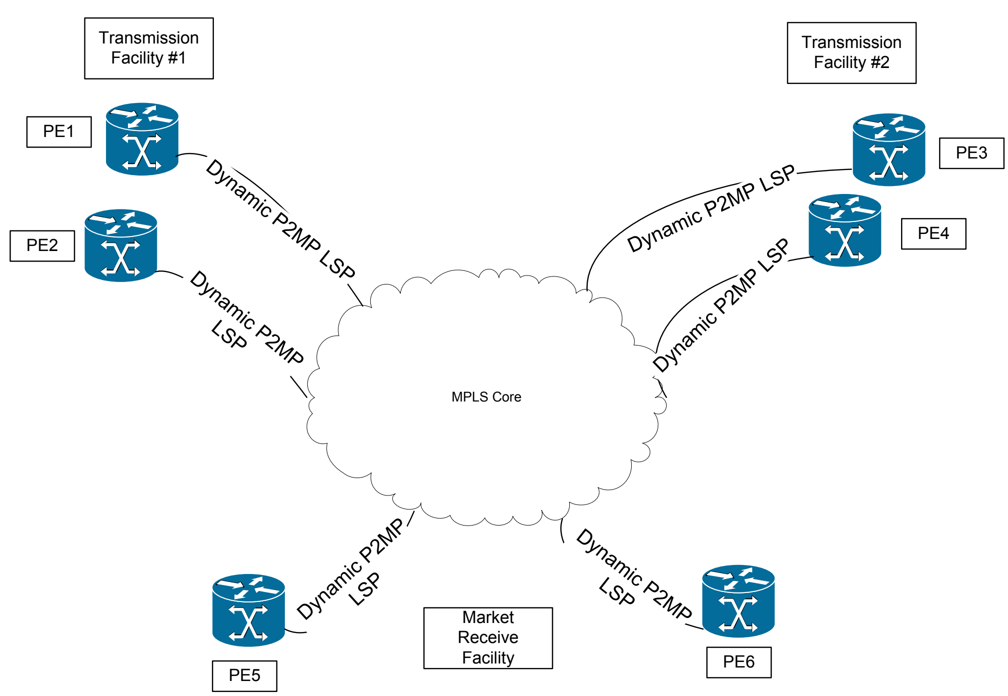

Figure 1—V2B 1.0 Distribution Model

backbone. In 2007 Cox began construction of a nationwide DWDM fiber optic transport network and immediately began converting almost all intermarket traffic (including residential and business high speed internet, telephony, and video on demand services) from leased circuits to wavelengths on the Cox national long haul DWDM backbone. The DWDM backbone terminated directly in Cox-owned data centers at each market, eliminating the need for metro transport to connect to carriers at colocation facilities for circuit delivery.

At the same time, Cox also started the V2B 1.0 initiative which centralized acquisition of national broadcast video content at two "master" headends located on the East and West coasts. As before, much of the content was acquired via satellite farms, but Cox's national transport backbone also enabled greater use of wireline (fiber) feeds acquired directly from content producers' studio locations around the country. In these cases the feeds would be converted to IP over Ethernet at the source location and ingested into the Layer 3 backbone network at the nearest backbone node location. From there the programming would be sent to the two East and West coast master headends. At the master headends, all the national content—whether acquired via satellite dish or directly from the studio—was fed into the Cox Layer 3 backbone network and distributed to the markets via layer 3 multicast. The multicast trees were set up individually for each programming station based on which markets were offering a given programming feed to their subscribers.

The two geo-redundant acquisition facilities would back each other up in case of a local failure (e.g. satellite heater failure or power failure), network isolation from one facility or another (e.g. due to multiple simultaneous fiber cuts) or content degradation. One of the facilities would serve as a primary and backup feed and the second facility would serve as a tertiary backup.

All content originating from the two facilities are IP packet based streams and during V2B 1.0 were delivered to markets via P2MP TE LSPs. Each transmission facility used two originating routers to protect against a node failure. Only one PE router at a given transmission facility transmitted the content at a given time. Dynamic failover mechanisms were in place to route the content to the back PE in case of a node failure. Each market received a copy of the content from Facility #1 and the same copy of the content from Facility #2 on each PE router. Figure 1 reflects a simplified version of the path that content delivery from each facility would take to get to a market. In general each transmission used diverse paths to every market, however in some instances the two paths would naturally converge during normal operating conditions. To ensure diversity for the P2MP LSPs permanent "coloring" of certain paths was employed to force desired routing patterns.

The benefits of this design included a 50 ms failover during network events (fiber cuts and hardware failures) due to the link protection mechanism built into the P2MP LSP technology. This technology pre-signals back-up paths, thereby bypassing failed links failed and building an alternative route to the destination irrespective of whether optical protection was enabled for the affected wavelength(s). It is important to note that in this design all V2B traffic was set at an Assured Forwarding Class of Service to ensure that should congestion occur, video traffic would flow unimpeded.

V2B 1.0 was a good solution for Cox because it allowed consolidation of channel acquisition efforts while also providing the ability to easily roll out new channels to every market. It also allowed deployment of video services on a platform and technology (MPLS) that operations teams were already comfortable troubleshooting.

V2B 1.0 Challenges

As long as the total bandwidth needed for the centrally acquired content stayed within the 2 Gbps range, using the Layer 3 IP backbone to distribute to the markets was the most efficient solution, given that most router-to-router links already consisted of several 10 Gbps connections. However, the explosion of HD content soon had bandwidth requirements far exceeding 2 Gbps. At this point Cox needed to reconsider the Layer 3 multicast approach which presented various challenges with respect to scalability. 6 Gbps of video traffic had been determined as the level where Layer 3 would no longer be cost effective as the transport mechanism, and in fact at 4 Gbps the video traffic started to become burdensome, presenting the following particular challenges:

QoS:

- What are the appropriate bandwidth thresholds across aggregates?
- What happens when a fiber cut causes transmission facility 1 and 2 to take the same path? Should one be protected over the other? The problem is that during a fiber cut or network failure it is possible for both facilities' content to be rerouted via the same bypass LSP. This causes the bit rate to double while the primary P2MP is being re-signaled. It also needlessly reduces throughput for lower QoS services such as high speed internet (or correspondingly raises costs to engineer significant extra failover capacity.) One solu-

tion to this problem is rigorous and labor intensive modeling and provisioning of the optical Shared Risk Link Groups (SRLG) in the router configurations.

Capacity Planning:

- With both transmission facilities sending a copy of the same content to each PE router, bandwidth scaling is prohibitive.
- At the time of deployment the video P2MP traffic was broken up into 5–8 LSPs to allow load balancing across aggregate links, i.e. to circumvent router link hashing algorithms that were based solely on LSP. However, during failure or congestion events all video bandwidth would still frequently be crowded onto a single link within a bundle. Overall capacity costs for both transport wavelengths and router interfaces and fabrics were kept high due to the escalating bandwidth needs for video, especially when failover capacity was considered.

V2B 2.0: Central Acquisition at Master Headends, Distribution via Layer 1 Multicast

V2B 2.0 was designed with the objectives of solving the QoS and capacity challenges without adding operational complexity. A solution was desired that would provide for at least 10 Gbps of video bandwidth to each market in a cost effective manner and it would need to be at least as reliable as the previous solutions. Ability to leverage the existing monitoring infrastructures was also desired, to avoid lengthy delays in the production network due to tools gaps. Finally, improved predictability of network behavior during outage events was also an objective.

The solution that was developed is based on broadcasting the superset of national content over dedicated wavelengths to each market. Two 10 Gbps wavelengths are employed, one that originates at an East Coast master headend, and the other on the West Coast. By setting the wavelengths up in a 'counter-propagating' tree-and-branch manner (Figure 2), each wavelength actually provides two diversely routed copies—dubbed 'A' and 'B', each representing one incoming side of the same wavelength—of the broadcast content from each master headend to each DWDM ROADM node in the market, i.e. four identical copies of the broadcast video content is provided to each market using the two wavelengths. The counter-propagating arrangement means that the wavelengths are used efficiently. The DWDM node acts as the initial switch point between the A and B feeds depending on network conditions, delivering a single East and West Coast feed to the local distribution network via the switch

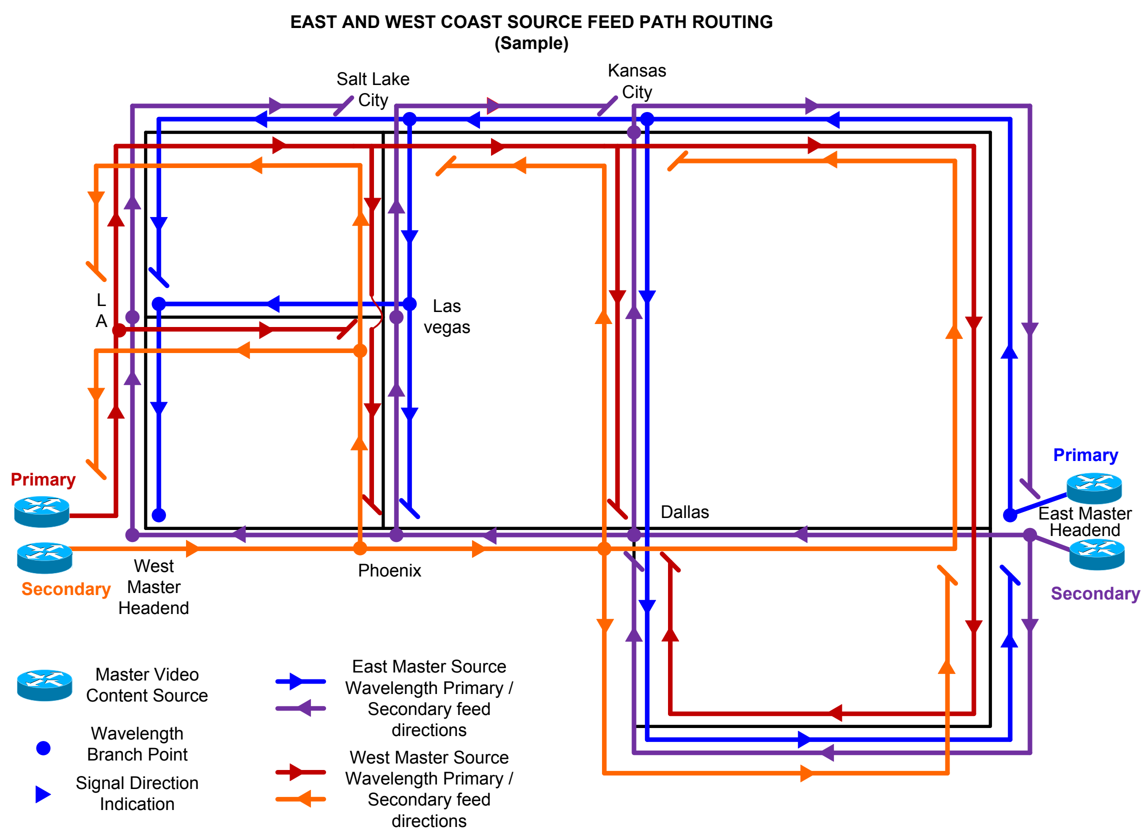

Figure 2—V2B 2.0 General Wavelength Multicast Tree and Branch Topology

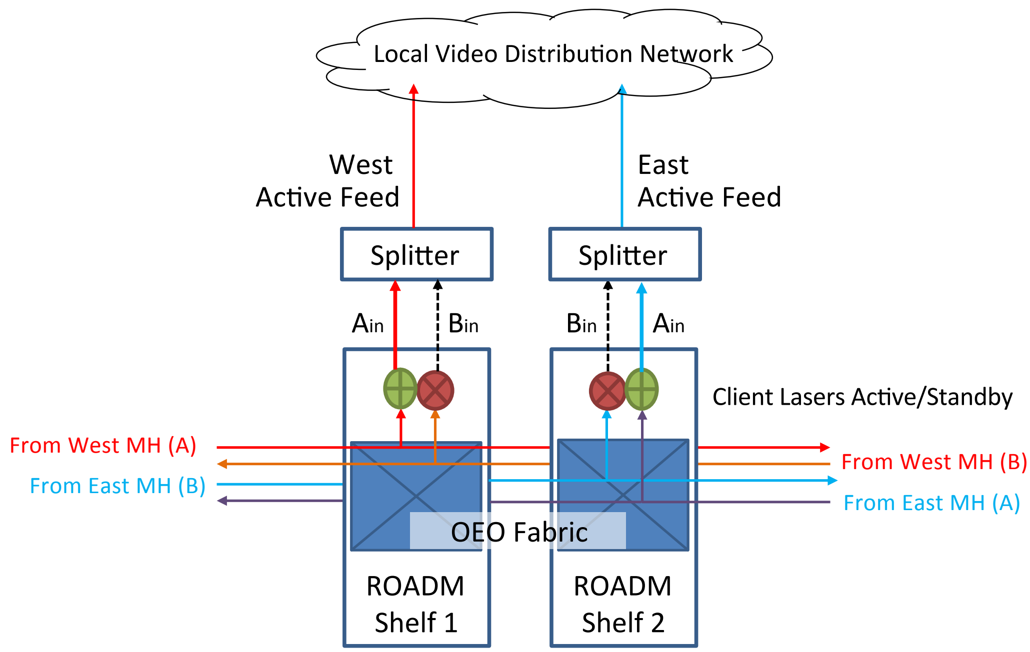

Figure 3—V2B 2.0 Typical Market DWDM Transport Detail

and splitter arrangement shown in Figure 3. The local distribution network provides the switch point between the East and West Coast feeds to handle conditions when the feed from the usual headend goes down altogether or experiences degradation. The selection of market-specific broadcast content (e.g. specific regional versions of certain programming such as sports-oriented networks) from within the superset of national programming channels also occurs locally prior to final distribution to subscribers over the access network, as well as assembly of local advertising zones and ad insertion, etc.

The V2B 2.0 solution offers the following reliability measures:

- Shelf-level diversity against transport-related failures. The drop points of the East and West Coast copies were engineered to land on separate ROADM shelves (Figure 3). This was made easier by the fact that the DWDM backbone network had been built initially with at least two shelves at each market node, with a consistent arrangement of wavelength bands on those initial shelves. Therefore, there were already shelves and cards in place to terminate two separate wavelengths (i.e. λ1 for the East feeds, λ1 for the West coast feeds.)
- Resiliency to at least two fiber cuts. The tree-and-branch designs of the wavelength multicasts were set up so that at least one of the four feeds will stay up on every fiber section even when there are cuts to two different fiber sections on the network. This was a challenge provided to Engineering by the Operations team during the initial design phases to ensure that video maintain maximum availability in the absence of a higher layer re-routing mechanism. While it was well understood that the resiliency to dual fiber cuts should be topologically possible, it was laborious to design manually and therefore expedited the creation of a tool to identify impacts of maintenance or fiber cuts, which was then used to create the final tree-and-branch design.
- Dual-level physical protection. The primary protection mechanism in the markets is to switch between the A and B feeds from the same master headend, which occurs transparently to the distribution network within a 50 ms timeframe. This switching is initiated mainly by a complete loss of signal or traffic due to fiber or transport equipment failures. The next level of switching is from one master headend source to the other, which is done by the distribution network and can take longer to occur. However, this switch can be initiated by a variety of other situations such as quality degradations due to video component failures, weather events or other headend issues.

V2B 2.0 Optical Design Details and Features

The optical design for V2B 2.0 was implemented fairly rapidly due to the initial dual-shelf layout established when the DWDM backbone was launched, i.e. the previously described consistent installation of dual-shelf node layouts with common wavelength bands at each site. Specifically, the ITU laser design utilized single-card 10-channel PIC-based laser banks, so the needed ITU lasers were already installed and operational even before V2B 2.0 was conceived. Additionally, an OEO design separates the ITU DWDM line optics from the client interfaces through the backplane switch fabrics allowing for a very elegant digital broadcast option, including support for sub-wavelength digital multicasts at rates of 1 Gbps, 2.5 Gbps, etc. The design enables arbitrary numbers of wavelength branch points to both clients and other OSP fiber sections to be constructed merely via cross-connects at each node without concern as to optical loss, OSNR levels, etc. The only equipment installations needed on the transport network to put V2B 2.0 into service were the client-facing optics and optical splitters. The only purpose of the optical splitters is to serve as a single-client interface to the switched A / B feed from each master headend (as well as a signal monitoring point).

Looking even further beyond the initial 10 Gbps wavelength, the cross-connects and wavelength reservations have already been provisioned for a second wavelength when the need arises for an additional 10 Gbps of capacity. At that time the only installations needed will again be the client optics and optical splitters. This frees the engineering staff from any planning efforts needed to account for changing broadcast channel lineups.

Additionally, the OEO tree-and-branch methodology allows branch points to be added or pruned with minimal service impact. Again, they can be performed simply by making necessary cross-connects without regard to any optical design. In the rare event of operator error during installation or cross-connect work, the reliability measures built into the design ensure that there will not be a complete video loss to any market as a result.

Other V2B 2.0 Advantages

The V2B 2.0 design provides benefits to other organizations beyond the groups associated with planning and designing the optical transport network.

Scalability

As mentioned previously, in the Layer 3 multicast model, insertion of new broadcast programming into the lineup required careful study of available capacity and impact on other non-video services during failover, as well as construction of new multicast trees to steer the new content to the correct market subset. In the V2B 2.0 model, adding a new programming channel is very straightforward; it is simply multiplexed with the existing channels into the master feeds, and then selected at the individual market level for further local distribution to subscribers. It frees the router and transport engineering groups from involvement in changes to the broadcast video content.

Network Operations Efficiencies and Improved Maintenance Practices

V2B 2.0 has reduced the operational overhead due to eliminating the network routing teams from troubleshooting. In V2B 1.0, issues affecting video service could be due to headend, router, or transport/fiber problems. The router layer has effectively been removed, reducing “finger-pointing” between technology groups and escalating problem resolution time. The complexity of the technology has been reduced due to the nature of the design, particularly in terms of the behavior during outage events. In the V2B 1.0 Layer 3 multicast model, problems were often due to complicated rerouting results due to network issues. Packet rerouting was difficult to predict and difficult to discern during an outage or maintenance event. With V2B 2.0, the relationship between fiber and equipment failures and the health of any given video feed at the markets is well understood, documented, and easily modeled using a spreadsheet tool developed for this purpose.

The failure analysis tool allows a user to indicate any fiber link as up or down and observe which feeds will remain healthy in every market. This allows for very precise notifications to market personnel in advance of a planned maintenance activity and also serves as a validation of the observed network status during an unplanned outage. The tool is not based on a static matrix of relationships, but on a dynamic model of the network as new branch points are added or removed. This is accomplished via a custom visual basic formula to calculate routability between endpoints using the Dijkstra shortest path routing algorithm. A sample readout of the tool is provided in Figure 4.

V2B 2.0 Challenges

A few challenges remain even with the V2B 2.0 model in place. While some of these challenges represent more of an inconvenience offset by improved reliability measures, others require maintaining some (reduced) level of continued operational support for a V2B 1.0-style model of delivery for certain markets and programming. This ongoing support though now has a much more manageable task of troubleshooting issues on the Layer 3 network, as issues will be related to specific programming channels instead of the entire broadcast lineup, or to specific markets instead of some possibly large number of markets being impacted simultaneously. These challenges include:

Input: Indicate Fiber link(s) Status:

Fiber Link	Up / Down
Atlanta to Dallas	Up
Atlanta to Macon	Up
Atlanta to Raleigh	Up
Baton Rouge to Dallas	Down
Baton Rouge to New Orleans	Up
Chandler to Dallas	Up
*	*
*	*
*	*

Results: V2B Feed Status by Market

Market / RDC	East Coast Primary	East Coast Secondary	West Coast Primary	West Coast Secondary	# Feeds Remaining Up	# Feeds Normally Up	# Feeds Lost
Baton Rouge	Down	Up	Up	Down	2	4	2
Chandler	Up	Up	N/A	N/A	2	2	0
Cleveland	Up	Up	Up	Up	4	4	0
El Cajon	Up	Up	N/A	N/A	2	2	0
Federal (Future)	N/A	N/A	N/A	N/A	N/A	N/A	N/A
*	*	*	*	*	*	*	*
*	*	*	*	*	*	*	*
*	*	*	*	*	*	*	*

Figure 4—V2B 2.0 Feed Status Assessment Tool, Sample Display

- **Regional content and broadcasting restrictions increase bandwidth usage on V2B wavelengths.** For many national networks, multiple versions need to be distributed to account for time zone differences (e.g. East and West time-shifted versions of a national network) or to allow for different targeted programming to be offered in an area (such as in sports-oriented networks.) The latter is of greater consequence for V2B 2.0 than the former, especially since the various regional versions of the programming are often oriented both towards providing geographically relevant viewing choices as well as abiding by broadcasting restrictions of content owners (e.g. Major League Baseball.) With V2B 2.0 this means many regional versions of programming are received at markets unnecessarily.
- **Exchanging local content between markets is inefficient on V2B 2.0.** Often programming produced locally within a Cox market is of interest to subscribers in specific nearby markets. To transport this content between markets on V2B 2.0 requires that it first be sent to the master headends for insertion into the national superset (over the routed network), only to be returned back to markets near the same location as it started. This uses up bandwidth both on the V2B 2.0 wavelengths and on the Layer 3 network. In these cases the more efficient solution is to send this type of content between the related markets on ad-hoc tunnels over the Layer 3 network and constrain the bandwidth usage to a much smaller portion of the network.
- **No return path via the V2B 2.0 Wavelengths.** Due to the counter-propagating design of the V2B 2.0 wavelengths, there is no way to transport a feed back from a market for centralized monitoring when there are unexplained quality issues etc. as both sides of the wavelength are carrying traffic downstream in a "ships passing in the night" fashion. While a feed can be extended around the network so that it terminates back in the originating master headend, the physical mesh topology of the network means that there will still be network sections on different branch paths that are different from what the master headend is receiving.
- **The "reconvergence" dilemma with 100 Gbps wavelengths.** Placing the national video traffic onto dedicated wavelengths represented a "deconvergence" of the video traffic at national backbone level, i.e. treating the video as a distinctly different, segregated service from the Layer 3 IP architecture used for all other services. This made sense in part because the amount of bandwidth needed

for broadcast video traffic aligned nicely with 10 Gbps increments that are the standard wavelength and router interface capacities currently in use. However, 100 Gbps wavelengths and (Ethernet) router interfaces lie ahead in the not too distant future, along with increased convergence between the router and optical control planes for Shared Risk Link Group awareness etc. At first glance this prospect seems to reopen the discussion of whether video traffic should be merged with other services over the routed network, under the assumption that the national broadcast video content is not likely to approach more than 20 or 30 Gbps for the foreseeable future, which would seem to remove the capacity management issues identified previously that were a major impetus for the transition from V2B 1.0 to V2B 2.0. The dilemma specifically would seem to be whether it makes sense to continue using valuable optical spectrum to deliver 10 Gbps wavelengths when 100 Gbps wavelengths begin to overtake the transport design. In practice, there will still be many smaller markets served primarily by 10 Gigabit Ethernet interfaces for some time, even if carried over 100 Gbps wavelengths. Additionally, OTN switching is becoming an almost standard feature of most long haul platforms, which will allow for the service-specific traffic to be transported efficiently without necessarily reaching the Layer 3 routed network supporting the other services. Thus, the benefits of V2B 2.0 will continue to be realized in the world of 100 Gigabit wavelengths and router interfaces regardless of whether the bandwidth needed approaches 100 Gbps.

- The "deconvergence" dilemma in the local markets. If deconverging the video onto dedicated wavelengths on the backbone network makes sense, an obvious question is whether it also makes sense in the markets. The answer (at this point in time) is still "no". One significant reason is that the subset of content needed within a market is still significantly less than 10 Gbps and even when combined with local broadcast content does not necessarily represent a majority of the video-based bandwidth being distributed. Other reasons are related to the fact that most markets have several different advertising zones that share a combination of dedicated and overlapping network facilities (transport, routing, access cable plant) that would make a deconverged model very inefficient and costly, requiring many instances of additional equipment within each market to reassemble zone-specific channel lineups from within the wavelength carrying the combined zones. Further, unlike the backbone network which is built on a single platform well-suited to scalable Layer 1 multicast, in the markets the optical transport networks frequently involves multiple platforms selected and sized for local requirements, having varying degrees or methods of support for Layer 1 multicast. Altogether these factors actually raise the costs and operational complexity of a deconverged V2B 2.0 model at the metro level compared with the current converged model in place in the markets.

Conclusion

The Layer 1 digital multicast approach to distributing national broadcast video content has allowed Cox to make optimal use of the DWDM long haul network. The design has offloaded large volumes of bandwidth from the Layer 3 network, reducing its costs and operational complexity while enabling a host of advantages in the distribution of the linear broadcast video, including rapid addition of new programming options, increased survivability, greater predictability, and improved maintenance and troubleshooting capabilities. Moving forward, as the routed network takes on increasing volumes of switched video content for over the top and interactive applications, the V2B 2.0 design will minimize the bandwidth duplicated for the same content while remaining compatible with next generation transport platforms and backbone architectures.

Acronyms and Abbreviations

DAC—Digital Access Control
DNCS—Digital Network Control System
DWDM—Dense Wave Division Multiplexing
HD—High Definition (Television)
LSP—(MPLS) Label Switched Path
OEO—Optical-to-Electrical-to- Optical
P2MP—Point to Multipoint
PIC—Photonic Integrated Circuit
QoS—Quality of Service
ROADM—Reconfigurable Optical Add-Drop Multiplexer
TE—Traffic Engineered
V2B—"Video to the Backbone"

This paper was originally published in the 2011 SCTE Cable-Tec Expo technical proceedings and was presented at the Expo on November 14 and 16, 2011.

Global Headquarters
169 Java Drive
Sunnyvale, CA 94089
USA
Tel: 1 408 572 5200
Fax: 1 408 572 5454
www.infinera.com

US Sales Contacts
sales-am@infinera.com

Asia and Pacific Rim
Infinera Asia Limited
391B Orchard Road
#23-01 Ngee Ann City
Tower B
Singapore 238874
Tel: 65 6832 8099
sales-apac@infinera.com

Europe, Middle East, Africa
City Point
1 Ropemaker Street London, EC2Y 9HT
UK
Tel: 44 207 153 1086
sales-emea@infinera.com

Customer Service and Technical Support
North America
Tel: 877 INF 5288
Outside North America
Tel: 1 408 572 5288
techsupport@infinera.com

infinera®