BEYOND BANDWIDTH MANAGEMENT: BUSINESS BENEFITS OF SWITCHED DIGITAL VIDEO IN CABLE

Lorenzo Bombelli
Director, Product Strategy and Management
Subscriber Network Systems
Scientific Atlanta, A Cisco Company

Abstract

Fueled by the advent of an open IP-based architecture that has catalyzed the development of reliable, cost-effective, and scalable solutions, Switched Digital Video (SDV) technology promises to fundamentally change how digital video is delivered over cable networks, enabling MSOs to offer consumers a wider variety of programming while effectively managing HFC network bandwidth.

Although bandwidth management is the primary driver for SDV in cable, from a business perspective SDV offers MSOs a number of additional benefits that reach beyond its pivotal role as a bandwidth management tool.

Unlike its telco cousin, the cable version of SDV is designed to operate over existing HFC infrastructure and to enable delivery of switched video services on the existing installed base of some 40 million MPEG set-tops that cannot decode a DOCSIS or IP stream, giving MSOs significant scale advantage in the introduction of new video services. Key switching features, like the ability to share QAM bandwidth between multiple services on a per-stream basis on a single QAM carrier give MSOs the same granularity and flexibility of service delivery to existing MPEG set-top boxes as any competitive service.

In contrast to traditional storage-based video on demand (VOD) technology, SDV is transparent to consumers. By preserving the consumer experience, SDV immediately helps MSOs give consumers more of what they want when they want it without changing the way they get it.

Because SDV fundamentally changes the model for bandwidth consumption from a linear model based on the program offering to one based on program viewership, SDV helps MSOs offer an extensive lineup of niche content ranging from local and other premium sports packages to ethnic programming, thus leveraging the "long tail" phenomenon to improve customer satisfaction, reduce churn, and generate new revenue streams from premium tiers.

A powerful byproduct of SDV systems is that they generate detailed viewership data that give MSOs unprecedented and direct visibility to consumer viewing behaviors, while maintaining the privacy of individual customers. This data can be used to improve program offerings, maximize the return on investment for each program offered, and drive growth in ad revenue by targeting ads more granularly.

Ultimately, switched digital video networks can be configured for switched unicast, enabling MSOs to deliver an individual copy of broadcast content to each consumer complete with highly targeted advertising and to offer features that enhance the viewing experience such as faster channel changes.

Finally, the new IP-based SDV architecture offers MSOs the opportunity to extend the reach of switched digital services to address the explosion of IP-based devices that are capturing an increasing share of video consumption.

<u>OPEN SWITCHED DIGITAL VIDEO ARCHITECTURE</u>

Over the past year, industry leaders have collaborated to develop an open architecture for delivery of SDV to the existing installed base of MPEG-based set-top boxes as highlighted in Figure 1.

Some of the key components and features of the SDV architecture include:

- An **SDV Client** protocol that can be natively integrated into set-top program guides and digital navigators from multiple vendors.

- A **Master Session and Resource Manager (SRM)** that is able to control and arbitrate bandwidth for multiple applications. Initially the master SRM manages the dynamic allocation of bandwidth between switched digital video and existing video on demand services. However, the architecture also allows for additional application servers to share system bandwidth.

- A **Switched Digital Video (SDV) Server**, a distributed SRM that is client-aware and constitutes the brains of the switched digital video network. The SDV server is designed to handle a large volume of real-time channel change and other service requests, and for high system availability and redundancy. It keeps track of all user requests and can use the information to improve system performance or to report viewership data and other key system parameters. The switched digital video server can be distributed or centralized, and is built on standards-based server hardware (e.g. Dell) and operating systems (e.g. Linux).

- **Session-Based QAM Modulators** that support IP Multicast (IGMPv3) and interoperability with standard IP switch/switch-routers. The session-based

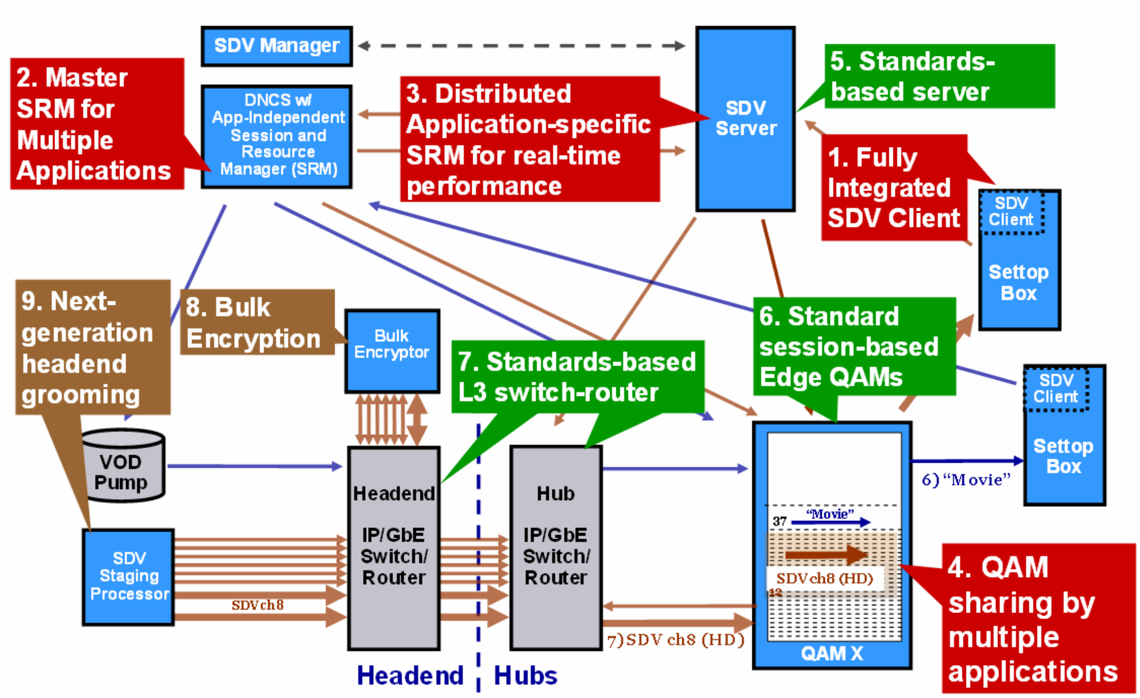

Figure 1. Key Components and Features of the Open Switched Digital Video Architecture

QAMs interface with both the Master SRM and SDV server using open, published interfaces to dynamically deliver switched video services over the HFC network, and enable QAM sharing between SDV and existing video on demand services.

- **Standards-Based Switch-Routers** that provide content distribution over local and regional fiber backbones. The key requirement that the SDV architecture places on the IP-based content distribution network is that it supports IP Multicast (IGMPv3) with Source-Specific Multicast (SSM) in order to provide source-level redundancy.

- Network-attached **Bulk Encryptors** that allow for independent scaling of encryption and edge QAM resources. Like the SDV server, the network-attached bulk encryptor is connected to the IP network over Ethernet (the bulk encryptor has multiple bi-directional Gigabit Ethernet ports) and can be centralized or distributed to achieve optimized cost and performance.

- A network-attached **Staging Processor** for grooming and processing of switched video streams. Initially, the role of the staging processor is to clamp the video streams to a known maximum bit rate to enable deterministic bandwidth allocation in the switched network. However, in the SDV architecture the staging processor is a network-attached appliance that may be able to perform additional functions such as ad insertion, rate shaping, transcoding between different digital video formats and alignment of key MPEG video frames (I-Frames). Although today the cost of MPEG processing is much higher than the cost of switching, and would burden the overall system if applied at the network edge, as higher density, lower cost video processing platforms are developed, the open SDV architecture allows for video processors to be applied seamlessly,- connected by IP across the network and scaled according to business requirements.

- **Standards-Based Reliability and Redundancy Mechanisms** that address both content and signaling networks to provide redundancy and service resiliency. These mechanisms protect against outages in sources (including encoders, staging processors and bulk encryptors), content distribution network, QAM modulators, and SDV servers. The SDV architecture also support service resiliency in the event of an HFC reverse path outage.

- An **SDV Manager** that features standard SNMP MIBs for automated configuration and management of system components from a single user interface.

- Architectural fit with future service convergence architectures such as Packet Cable Multimedia and IP Multimedia Subsystem Architecture.

The referenced white paper, *"An Open Architecture for Switched Digital Services in HFC Networks"* [1], provides a detailed overview of the key elements and operation of a switched digital video system using this open architecture.

The advent of an open SDV architecture has accelerated the development of technologies and solutions for delivery of SDV services over HFC networks. It is more reliable, more cost-effective and more scalable than existing MPEG-based architectures. As such, to MSOs, a switched digital video system is more than just a powerful bandwidth management tool. It is a key stepping stone in the evolution to next generation networks that not only pays for itself by delivering an immediate return on investment based on bandwidth management, but also delivers a number of key benefits that add business value to MSOs both immediately and going forward.

ENABLING EXISTING MPEG SET-TOP BOXES

One of the most important features of the open SDV architecture is that it enables switched delivery of video services over IP network infrastructure to the existing installed base of more than 40 million MPEG set-top boxes. The alternative would be to overlay switched services to IP-based set-tops. An IP overlay would not only fail to deliver immediate bandwidth savings (additional bandwidth would be required to support the overlay), but also to address the delivery of new services to existing digital subscribers without replacing their set-top box and incurring significant capital, operational and opportunity cost.

Interestingly, because the open, IP-based SDV architecture reduces the bandwidth required to deliver both existing and new services to existing MPEG set-tops, it also catalyzes the evolution to all digital and end-to-end IP-based networks by freeing up bandwidth for services that require a video overlay such as digital simulcast and video over DOCSIS.

QAM BANDWIDTH SHARING

As shown in Figure 2 below, the open SDV architecture is designed to enable the coexistence of SDV and existing VOD services, and enables dynamic bandwidth sharing between multiple services sharing the same edge resources.

An application independent SRM enables sharing of edge resources (QAM modulators for an HFC system) among the various interactive applications requiring HFC bandwidth. Such sharing of QAM resources is normally considered to occur in one of two ways: inter-carrier sharing, and intra-carrier sharing.

Inter-carrier QAM sharing, the simplest and most straightforward method of sharing QAM resources among applications, involves the assignment of individual RF carriers to specific applications. Some of the RF carriers within a QAM chassis, for example, may be assigned to VOD while others may be assigned to SDV. Inter-carrier QAM sharing is less dynamic, but is simpler to implement and does not require the intelligence to arbitrate between more than one application demanding share of a given carrier.

Intra-carrier QAM sharing is more complex but more powerful than inter-carrier QAM sharing. It relies on some level of intelligence to dynamically make decisions on the priority of two or more applications requesting access to stream capacity within a QAM carrier. To accomplish this, carriers are designated on the SRM as sharable among more than one application. It is then up to the SRM to assign bandwidth on each carrier upon request based on some set of rules or business intelligence. It may be as simple as prioritizing the applications as a whole, or as complicated as prioritizing individual streams based on the value of content, time of day or any of a number of factors. The requirement for decision-making intelligence, intra-carrier QAM sharing calls for QAM modulators to be entirely session-based.

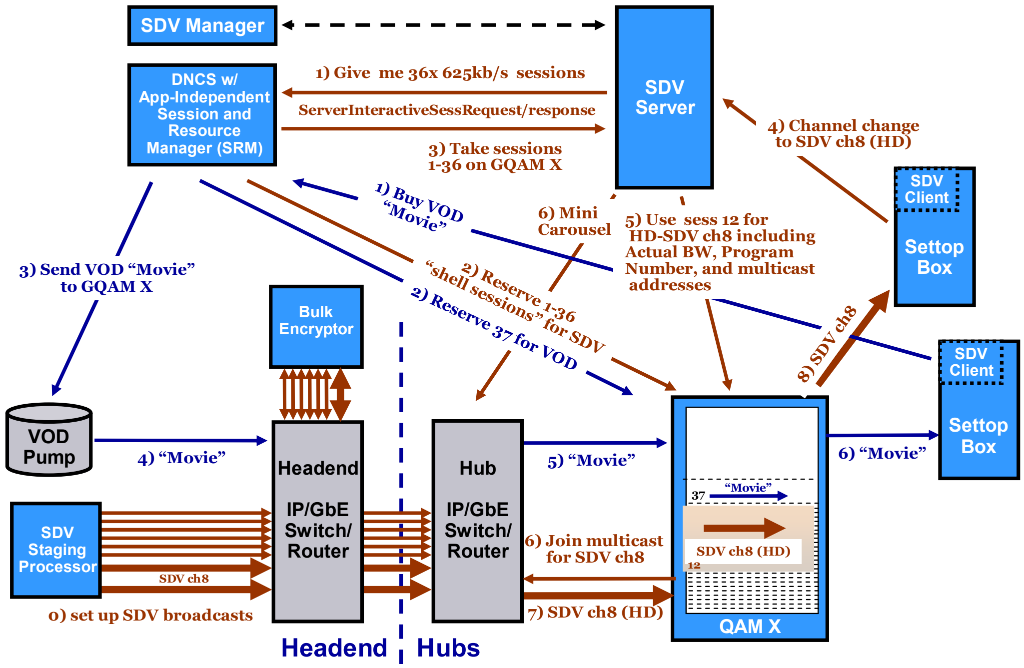

Figure 2. SDV and VOD Coexistence

TRANSPARENCY, CUSTOMIZATION AND "LONG TAIL" CONTENT

Switched digital video extends the capability of MSO networks to enable the delivery of content on demand without changing the average consumer's "lean back" viewing experience.

In order to view video on demand on a traditional storage based VOD system, a consumer searches for a given piece of content, selects it, and then waits approximately 5-10 seconds for the selected content stream to appear on the screen. The VOD experience is very different than typical TV viewing, wherein a consumer tunes to a favorite channel or surfs through channels to see "what's on". He then might stay tuned to that channel as long as it continues to entertain him.

Because SDV is designed to behave, from the consumer's perspective, exactly like broadcast TV, it must deliver the same experience as regular TV viewing. However, because it takes advantage of viewing statistics to deliver video streams only if they are being watched, it enables the delivery of many more channels, allowing MSOs to offer content that provides high value to small groups of users, sometimes referred to as "long tail" content [11].

Opportunities to deliver such programming, ranging from local and other premium sports tiers to international and ethnic programs (as illustrated in Figure 3), is now offered by satellite and can drive growth in new subscribers and improvement in customer retention if offered as part of an MSO's bundle of services.

One of the challenges in delivering these channels over cable networks is that there are pockets of demand for many of these channels in most cities, and certainly in any metropolitan area. As a result, in order to

Figure 3. Examples of International Programming Options

meet the needs of the entire subscriber base, an MSO would have to offer many or even all of the packages shown above. In the traditional linear broadcast delivery model this plan would be impossible. With switched digital video—which capitalizes on geographic diversity, viewing statistics and the efficiency of multicast—however, most MSO networks can deliver most (if not all) of the programs shown above.

VIEWERSHIP DATA AND CLIENT AWARENESS – GETTING CLOSER TO CONSUMERS

A valuable byproduct of switched digital video systems is very detailed viewership data. The SDV system has visibility to every channel change request in real-time, whether it is for a switched, broadcast digital, or analog channel. It also receives and stores non-real-time reports from set-tops on all other remote control activity. Although the viewership and user activity data is primarily used to optimize system performance, the data can be reported from the SDV server in a secure manner that protects consumer privacy.

The enhanced visibility to consumer behavior can be used to continually fine tune and target service offerings, improve marketing campaigns, and increasingly target advertising to drive ad revenue growth.

The SDV server is also aware of the capabilities of each SDV client (set-top or other device) that it serves. This client-awareness combined with rich viewership and activity data can be used for a number of key operational benefits including facilitating the rollout of new set-tops or devices that support advanced video formats, or delivering different streams to different devices depending on the bandwidth or format capabilities of each device.

Combined with powerful network-attached video processing and digital ad insertion, these capabilities may prove to be key drivers of revenue growth for MSOs.

SWITCHED UNICAST

Taken to the limit, targeted advertising leads ultimately to a switched unicast model where an individual copy of each program is delivered to each customer even if the program is viewed by more than one customer.

As with any capability, switched unicast comes at a cost. First, switched unicast by its very nature forfeits much of the efficiency offered by switched multicast, and, as a result, fundamentally costs more than switched multicast. Second, the cost of the video processing technology required for digital program (ad) insertion (DPI) today is significantly higher than the cost of switched video, and makes it prohibitively expensive to deliver individualized ads.

The business model for highly targeted TV ads has yet to be proven and may not become viable for some time. However, as switched video systems give MSOs greater visibility to consumer demand, as video processing technology continues to improve at Moore's Law rates, and as competition for consumer attention intensifies, it is not a stretch to envision a day when highly targeted, if not unicast, streams become a business reality.

Given the current cost of video processing technology and the lack of a proven business model, it may not make sense to broadly deploy switched unicast capability initially. Nevertheless, an open, IP-based SDV architecture allows MSOs to introduce higher performance, more cost-effective, higher density network-attached, next-generation video processors to switched video networks to deliver increasingly customized content.

FAST CHANNEL CHANGE

The SDV architecture offers the opportunity to significantly improve channel change times. There are a number of methods available for speeding up channel change times. One method that is gaining momentum involves delivering a switched unicast stream to each subscriber and starting each stream with an I-Frame, an MPEG frame that can immediately resolved by the receiver. This approach, combined with other techniques, can significantly reduce channel change times and make them faster than they are in digital broadcast networks today.

Like individualized advertising, fast channel change is a powerful tool that is enabled by SDV and is available to MSOs to be used as required by business and competitive needs. Similarly, fast channel change will benefit from performance, cost and density improvements in video processing, and may become more viable with the advent of next-generation network-attached video processors.

A key additional consideration in evaluating the fast channel change feature is that the fast channel change capability will only be available on channels in the switched video tier. As a result, consumers may become less satisfied with their experience in viewing programs that are left in the traditional linear delivery tiers. Preservation of a consistent consumer experience is a key consideration in rolling out any new service or feature. In the case of fast channel change, it is likely to be driven, as are many new features in cable, by competition, and may accelerate a wholesale transition to switching all digital video services in order to preserve a consistent consumer experience.

EVOLUTION TO NEXT GENERATION ARCHITECTURES

As we described above and shown in Figure 4, the SDV architecture described here is an important element of a comprehensive and evolutionary next generation architecture strategy.

First and foremost the SDV architecture described here provides a critical first step from today's linear broadcast delivery model to a non-linear switched delivery model based on IP core infrastructure, and acts as a catalyst for the evolution to all digital and end-to-end IP-based services by freeing up bandwidth in the network.

In addition, because the open SDV architecture is based on end-to-end IP infrastructure using open protocols and interfaces, it enables the delivery of multicast video streams from the content source all the all the way to the IP network edge, where session-based edge QAM modulators output MPEG video over RF to existing digital set-tops.

As such, it offers MSOs the option to distribute existing and new services including video on demand, digital simulcast, and DOCSIS with switched digital video on a converged IP network infrastructure (whether it be converged on a single or multiple wavelengths is up to the MSO).

It also provides an immediate business driver (bandwidth management) to extend the end-point of IP convergence beyond the headend. Traditionally video has been broadcast over 1310 or 1550 nm fiber as an overlay to the IP data backbone. The SDV architecture extends the IP network to the hub where IP services are launched today

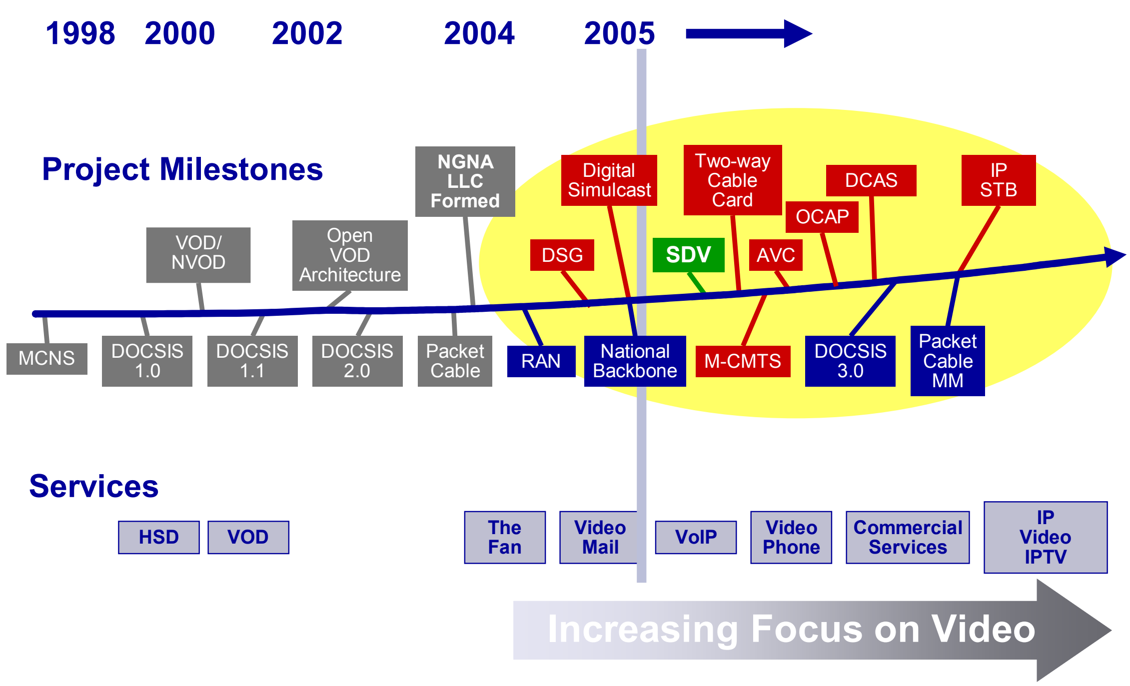

Figure 4. Evolution to Next Generation Architectures

over CMTS, an important step toward end-to-end IP convergence and next generation architecture.

Because one of the main goals of the SDV architecture is to address the existing installed base of MPEG set-tops, it features an out-of-band signaling mechanism that is structurally consistent with the Packet Cable Multimedia (PCMM) and IP Multimedia Subsystem (IMS) architecture models for non-IP clients. However, the SDV architecture can also be evolved to support IP-based clients and, in conjunction with the DOCSIS, PCMM and IMS architecture models, to provide a true multi-service, multi-client architecture that enables any-source to any-user IP-based video delivery.

SUMMARY

In summary, although bandwidth management is the primary driver for switched digital video in cable, the open, IP-based SDV architecture that is being embraced by MSOs offers a number of additional benefits and capabilities, including:

- The delivery of switched services to existing MPEG digital set-tops
- QAM bandwidth sharing between SDV, VOD and future on demand services
- A transparent consumer experience compared with traditional digital broadcast cable TV
- The ability to launch new premium and niche programming tiers
- Rich viewership data to help get closer to customers and better understand their needs
- A business-driven path toward increasing customization of content and advertising
- Powerful competitive features such as fast channel change
- A sure-footed evolutionary path to IP-based next-generation architectures

Thus, switched digital video not only delivers immediate benefits in the form of spectrum savings and increased efficiency, but it also presents great opportunities for sources of competitive advantage going forward.

REFERENCES

1. An Open Architecture for Switched Digital Services in HEF Networks, Luis Rovira, Lorenzo Bombelli, Paul Brooks, January 2006
2. Switched Multicast Services for HFC, Luis Rovira, October, 2005
3. Time Warner Cable Switched Broadcast Channel Change Message Interface Specification TWC-SDB-CCMIS-2_12, October 12, 2005
4. Time Warner Cable Switched Digital Broadcast Mini Carousel Message Interface Specification TWC-SDB-MINI CAROUSELMIS-2_1, October 12, 2005
5. SDB Server/RM – Edge Device Interface 1.2, September 9, 2005
6. Master SRM – Edge Device Interface 0.5, October 19, 2005
7. Session Setup Protocol Extension for Server Initiated Setup, Version 1.8 September 14, 2005

8. ISO 13818 part 6, Digital Storage Media Command and Control (DSM-CC)

9. TWC Pegasus Session Setup Protocol (SSP) Version 2.3

10. ISO 13818 part 1, MPEG Systems Specification

11. The Long Tail, Chris Anderson, Wired Magazine, October, 2004