White Paper

The Economics of Switched Digital Video

Gregory Hardy
Vice President of Business Development
SPVTG - Cisco

Benjamin Chen
Staff Applications Engineer
SPVTG – Cisco

Met Lebar
Transmission Systems Director
Time Warner Cable / Oceanic Division

Overview

Switched digital video (SDV) technologies are quickly gaining acceptance into the cable technology mainstream, partly because cable providers need to preserve bandwidth for popular-but-bulky services -- such as high-definition television (HDTV).

As rising numbers of high-definition programs drive video switches into more and more cable systems, it is important to be as accurate and inclusive as possible about the costs associated with this type of bandwidth reclamation. Those costs can be fixed and variable; they can be explicit and implicit.

The fixed costs of switching video are reasonably easy to determine. Its component parts, at a hardware level, include bulk encryptors, staging processors for rate clamping, switching servers, and quadrature amplitude modulators (QAMs).

In addition to that, though, are the implicit and variable cost implications of SDV. These costs are usually linked to preexisting topological conditions, and whether or not any further laser segmenting or node splitting is required. One such example is Video on Demand (VoD), and whether the launch of that service used centralized vs. distributed server placement. [1]

Another implicit and topology-dependent cost of SDV involves the total number of tuners[2] in a given network segment, where the network segment is defined by the number of 6-MHz digital channels allocated to the switch (the "QAM density"). In this document they are referred to as "tuner groups," not the more common VoD term "serving-area groups", because the number of tuners is a more accurate way to define the architecture-related economics of switching.

From a dollars and cents perspective, the cost calculations for SDV involve a strategic decision: How best to share resources across a grouping of households. Or, more specifically, how many QAM-modulated digital channels (resources) to apply to a given tuner group (users), as defined by the number of attached and simultaneously active tuners. A 1000-tuner group with 4 digital (QAM) channels, for instance, carries different cost implications than 8 QAMs deployed over a 500-tuner group.

In general, the most economic sweet-spot for smaller operators considering SDV is a design that spreads larger tuner groups across fewer digital channels -- fewer QAMs -–– meaning that more tuners, inside more digital set-top boxes, in more homes are requesting resources from a smaller number of QAM channel resources.

This configuration lets the smaller operator load the switch with an appreciable number of the least-watched standard and high-definition channels, while conserving the number of QAM channels allocated to the switch.

For such an economic assessment of switched digital video to be contextually accurate, though, it is also necessary to include the costs associated with "the rest of the story" after the sale, but before any appreciable bandwidth preservation accrues to the service provider.

Implicit costs are the items that do not necessarily factor into the equipment and software portion of an SDV deployment. For a service provider, these costs can add up, in a real, day-to-day, find-budget-for-it way: power and cooling, for example. In addition, there are costs related to optical network configuration or reconfiguration to get tuner groups and QAMs optimally aligned -- as well as all the other "make-ready" odds-and-ends that inevitably accompany an SDV deployment.

This paper aims to assist the "SDV-interested" readers to develop a reasoned and inclusive cost picture for adding SDV into their bandwidth preservation mix.

The Logic of Preserving Bandwidth by Using SDV

There are two primary ways to preserve bandwidth using SDV. The first is to increase the number of QAM channels per tuner group. Doing so increases the number of standard- or high-definition streams that you can apply to the switching pool -- but it assumes the existence of available RF carriers (empty QAMs).

The second way to preserve bandwidth via switching is to reduce the tuner-group size; fewer tuners sharing the same number of QAMs dedicated to switched video.

For that reason, this paper describes two such SDV deployments. Both use data from actual, working SDV systems. System 1, with 100,000 home passings and 36,000 digital cable customers, switches digital video content in 4 QAMs, serving 1,000 tuners each.

System 2 is larger. It passes 400,000 homes, serves 144,000 digital set-tops, and allocates 8 QAMs for switching to 500 tuners. Both systems assume 25 percent penetration of dual-tuner digital video recorders (DVRs), lifting the total tuner counts to 45,000 for System 1, and 180,000 for System 2.

---

[1] Centralized architectures can use cost-effective Gigabit Ethernet to move titles inexpensively from one server location into multiple distribution hubs; distributed topologies disperse multiple servers to the edges of the plant, keeping popular titles closer to consumers. The centralized approach is persisting in SDV architecture design, again because of the cost advantages of Gigabit Ethernet transport.

[2] A "tuner," in this sense, is described independently of a digital set-top box (STB) because an increasing number of digital STBs contain more than one tuner. Consider, for instance, the home that consumes television over three dual-tuner digital video recorders (DVRs). That home potentially populates narrowcast (on-demand and switched broadcast) television across six tuners, if both tuners within all three boxes are simultaneously displaying or recording digital TV content.

Table 1 summarizes the switch-related attributes of the two systems.

Table 1. Switch-Specific Characteristics of Systems 1 and 2

Switching Attributes	System 1	System 2
Homes passed	100,000	400,000
Digital set-top boxes served	36,000	144,000
DVR penetration	25%	25%
Total tuners served	45,000	180,000
Tuner-group size	1,000	500
Number of QAMs switched	4	8

Figure 1 illustrates the bandwidth savings achievable by the two systems. The x-axis shows the number of standard-definition video streams, and the y-axis shows the number of QAM carriers needed to carry those streams.

The diagonal black line in Figure 1 is the basis point: the number of QAMs required to carry linear, standard-definition programming, in traditional broadcast style, with no switching. In essence, it is the "now": carrying 150 standard-definition digital video streams requires roughly 13 QAMs, preswitching. (That number assumes MPEG-2 compression with a gain of approximately 20 percent from statistical multiplexing).

The solid, curved lines show the number of QAMs and tuners allocated to the video switch. The blue line characterizes the configuration of System 1, for instance, and the orange line represents System 2. The red line indicates what can happen with an even smaller tuner-group size - 250 tuners.

Subtracting the solid colored lines from the straight, diagonal line shows (in the dashed, colored lines) the number of QAMs saved by switching -- and therefore, the amount of preserved bandwidth. For instance, System 1 (in blue) is capable of saving 5 QAMs, whereas System 2 (in orange) can likely recover 8 QAMs.

If System 1 wanted to reclaim more QAMs past the inflection point (the intersection of dotted and dashed lines), next steps would involve resegmenting laser configurations or splitting nodes, to ultimately get to smaller tuner groupings. For System 1 to achieve the same bandwidth savings as System 2, for instance, it would need to go to a 500-tuner group; for System 2 to jump to a savings of 10 QAMs (the red line), it would need to reduce tuner-group size to 250.

Oversubscription Rates

The previous section illustrated how to think through potential bandwidth savings, based on the two key elements of SDV system architecture: QAM channel resources and tuner-group sizes. Now let's go deeper on what that means from the perspective of how many streams you can allocate to switching.

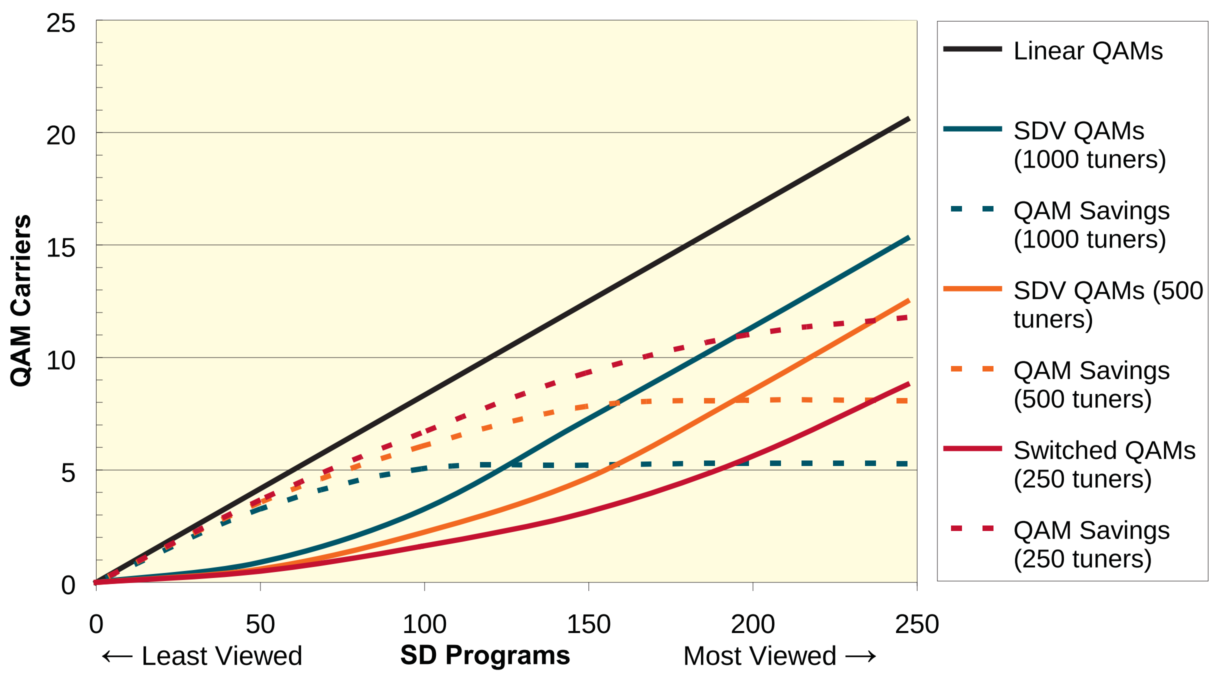

Figure 1. SDV Bandwidth Required and Bandwidth Saved for Various Tuner-Group Sizes

A common term in video switching is the "oversubscription rate". Oversubscription is the ratio between the sum total bandwidth of the switched programs and the bandwidth actually required to switch these programs. It is a way of drawing a line in the sand, so to speak, about how to harness the fact that most programs on most channels are not watched at any given time of day.

For example, carrying 100 standard-definition programs compressed to 3.75 Mbps, and 20 high-definition programs at 15 Mbps would normally require 675 Mbps of bandwidth, or approximately 18 QAM channels.

(100 x 3.75) + (20 x 15) = 675 Mbps ÷ 37.5 Mbps per QAM = 18 QAMs

If these same programs were to instead require only 6 QAMs, because of the bandwidth-preserving nature of SDV, the oversubscription ratio would be:

18 ÷ 6 = 3.0

Selecting an oversubscription ratio is highly dependent on three elements:

- Tuner-group size
- The relative popularity of the programs being switched
- The mix of standard- and high-definition programs: For extremely popular programming, this ratio can approach 1.0, meaning that oversubscription is not possible without blocking. These popular programs should remain on the broadcast tier and not be switched. For niche and so-called "long-tail" content, however, the ratio may be as high as 6 or 7. Most operators active with SDV are switching a blend of niche and semipopular programming to achieve oversubscription in the range of 2.5 to 4.1.

System 1

The smaller system, using 4 SDV QAMs with 1,000 tuners per tuner group, might decide to switch 60 standard-definition channels and 10 high-definition channels. The result would be a linear bandwidth requirement of 10 QAMs.

(60 x 3.75) + (10 x 15) = 375 Mbps ÷ 37.5 Mbps per QAM = 10 QAMs

By applying an oversubscription assumption of 2.5, the system 1 QAM requirements drop to 4:

10 ÷ 2.5 = 4 QAMs

System 1 has now freed 6 digital channel slots, which can be used for any number of new services. Perhaps they could be adding linear high-definition streams, for instance, or something completely different, such as adding the channel-bonding aspects of the DOCSIS® 3.0 cable modem service to achieve substantially faster downstream broadband service rates.

Note: Although System 1 could use these additional 6 channels to add 12 to 18 more linear high-definition channels, they offer an even greater return if used to switch even more "lightly viewed" standard- and high-definition program streams.

New standard- and high-definition video content is ideal for switching. Viewership on new channels is typically low, by default -- people are not yet aware of it. Allocating new and niche content to the switch allows the large or small cable operator to achieve very high oversubscription ratios. Subsequently, as the content becomes more popular, the operator can adjust the SDV lineup accordingly -- program streams can move in and out of the switched tier based on viewing popularity.

System 2

System 2 will achieve higher oversubscription than System 1 for the same channels because of the smaller tuner-group size of 500. The same SDV tier of 60 standard-definition streams and 10 high-definition streams might require only 3 QAMs, for instance -- resulting in an oversubscription of 3.3.

(60 x 3.75) + (10 x 15) = 375 Mbps ÷ 37.5 Mbps per QAM = 10 QAMs
10 ÷ 3.3 = 3 QAMs

By allocating 8 total QAMs for SDV, System 2 can begin adding to the switched tier some of the more popular channels in its current lineup (thereby lowering the overall oversubscription ratio), or it can launch substantially more lightly viewed standard- and high-definition programs.

Even a conservative over-subscription rate of 3.5 for new premium or niche content would result in capacity for over 40 new linear high-definition channels, within those remaining 5 QAMs.

The Fixed Costs of SDV

The same fixed startup costs apply to all system sizes considering SDV -- whether they are readying it for one set-top box or a million. Those costs are as follows:

- Headend components: Headend components include bulk encryptors, staging processors for rate clamping (for example, setting incoming, MPEG-2 compressed standard-definition video at a constant bit rate of 3.75 Mbps, or high-definition video at 15 Mbps), and software necessary to upgrade headend controllers (the Digital Network Control System [DNCS] in the case of Cisco® controllers.
- SDV servers: The "brains" of any SDV system, the servers are necessary to perform real-time switching and create different switched sessions. SDV server costs are variable, and organized to scale in step with system size. SDV servers typically represent a substantial portion of SDV deployment costs, and are highly dependent on the topology of the system.

- QAMs: Quadrature amplitude modulators are typically priced by density -- how many channels per chassis or per port, for instance. While QAM pricing is tracking swiftly downward, the devices are becoming increasingly "smart". Future QAMs will be required to support existing and forthcoming SDV protocols.
- Services and training: Typically, SDV deployments include a provision for "QAM rack-and-stack" services, including provisioning. Some operators want more; some want less.

The "Soft" Costs of SDV

Like any new technology -- bandwidth-saving or otherwise -- implementation always extends beyond the Bill of Materials. In the case of SDV, two primary variables can affect ongoing costs.

One is the assumptions made about tuner-group sizes, which tend to be highly variable. This number matters because more consumers are buying HDTVs than are signing up for digital TV service. This falls into the category of "a good problem to have," but it nonetheless impacts the timing of architectural decisions.

More subscribing HDTV homes necessarily lifts the raw number of digital tuners requiring QAM resources. In general, since 2005 cable operators have experienced 20 percent annual growth in the number of tuners they are required to support. This trend is significant, and it is unlikely to reverse itself, particularly as HDTV set prices continue to drop into mainstream affordability.

Rising consumer adoption of HDTV sets and services is particularly important for its impact on the *design* for SDV - the more tuners that appear in the network, the higher the possibility that the original design assumptions for tuner-group sizing could be exceeded sooner than anticipated.

The other variable is the addition of new channels. These days, most new channel launches happen in high definition, not standard definition. While although low-popularity and niche channels can be rapidly absorbed into existing tuner group sizes, any high-popularity and regularly watched channels can constrain available capacity, and should not be placed in the switching group.

Put another way, a heavily watched channel will appear in multiple tuner groups. More people watching, more neighborhoods watching, more active tuners per tuner group. If a channel is present in many or most tuner groups, it is unsuitable for switching. If it is niche, or lightly viewed, it is highly suitable for switching.

It is therefore the "semi-popular" or "occasionally popular" channels that run the risk of leading an operator into blocking conditions. These channels should be regularly monitored, within each tuner group, to baseline their viewing popularity and thus their suitability for switching.

If either variable exceeds the original design assumptions -- often made when VoD was first implemented -- the effect will be felt in both timing and expense. Reducing tuner-group sizes takes time; adding QAMs requires time, space, and budget.

In addition, narrowcast QAMs differ from broadcast QAMs because they are necessarily added into each distribution hub and feed a single tuner group -- not added into one, centralized location that services an entire system. That again translates into time, space, and budget.

Likewise for the necessary adjustments to the combining network, which are also complex and time-consuming.

Consider the system which operates out of 20 distribution hubs and needs to add QAMs. For SDV, those QAM additions will occur 20 times, one for each hub. Therefore, each hub must have the rack space, powering resources, and cooling to accommodate the extra QAMs. In some cases, this situation may even require enlarging the power feed to the building.

Likewise, when the current tuner-group size begins to exceed the design parameters, the next step is node splitting. Take System 1, with its 1,000 tuner-group sizing, as an example. If that system was designed for 1,000 tuners but built to 800 tuners, the 20-percent year-over-year tuner growth impact likely means that System 1 anticipated TG size will be exceeded within a year.

Realistically speaking, it is important to keep in mind that adding shelf space for high- and standard-definition programming via SDV is not operationally parallel to the known plant and bandwidth upgrades of the past. In other words, it is not a matter of just adding bandwidth. By contrast, the bandwidth afforded by SDV will require regular monitoring and adjustments -- at least annually -- to keep up with increased tuner and channel counts.

Adding It All Up

Based on actual cost data from a working SDV system, which has attributes similar to but not exactly the same as System 2, the all-in cost to implement SDV is in the range of $12 to $16 per home passed. Again, that's a system with turner group sizing of 500, over 8 QAMs. In this cost scenario, 40 percent of homes passed are digital cable subscribers, using an average of 1.5 digital set-top boxes per household. From a tuner demand standpoint, it is assumed that 25 percent of the installed footprint consists of dual-tuner models.

The costs break down as follows:

Vendor-supplied gear will cost between $6 and $7 per home passed, including a variety of components, most of which have already been discussed: QAMs, SDV servers, encryptors, and rate clampers. This cost assessment also includes vendor-supplied services for testing and provisioning the SDV gear, as well as bundled-in tools and annual support costs.

Plant segmentation, estimated at $2 to $3 per home passed, includes costs for laser resegmenting and node splitting, but is variable. Costs will vary depending on the extent to which tuner groups have already been split. It is important to note that the greatest potential for cost variations occurs on this line item. Often existing tuner groups are initially sized in the 1,000 to 2,000 (or larger) range based on the initial VoD deployment of the system. In this case a more significant and costly node segmentation effort will be required.

Hub labor and materials, important in systems with multiple distribution hubs, are expected to cost between $2 and $3 per home passed. Within that cost are items like racking and wiring services, provided by the vendor, as well as any upgrades necessary to the powering and HVAC systems. Also included in this figure is the cost to adjust RF combining networks.

IP network additions, necessary to establish overall IP connectivity, may be needed to transport more "long-tail" standard- and high-definition content over the IP network. For instance, consider the shift toward pulling in more user-generated video, from the Internet through to the TV. This process will likely require more router/switch ports, and cards to accommodate additional servers and QAMs. In this estimation, the cost for IP networking is $1 to $2 per home passed, with an additional $4 per home passed if new fiber is required.

Lastly, miscellaneous and project management costs are estimated at $1 per home passed.

All in, on the low side, that puts the cost of adding SDV as a bandwidth- expansion technique at $12 per home passed, and can go as high as $16 per home passed.

Conclusion

Gathering costs for SDV is not as simple as the Bill of Materials, because much of the calculations are based on variables that were set 'in stone' when the first narrowcast cable service was implemented -- VoD. For that reason, a working knowledge of how many unoccupied QAMs are available to be allocated to the switch, as well as existing tuner groupings, is recommended.

Still, though, for systems sharing the characteristics outlined in this paper - System 1, with 4 QAMs over 1,000 tuners, and System 2, with 8 QAMs over 500 tuners -- can achieve substantial bandwidth savings. System 1, in this calculation, allocated 60 standard-definition and 10 high-definition streams to its switching tier, applied a reasonable oversubscription assumption (2.5:1), and was able to reclaim 6 digital channels. Without switching, System 1 would have been required to allocate 10 full QAM channels to accommodate that loading. System 2, with the same switching tier components, applied a more aggressive oversubscription assumption (3.3), and reclaimed 7 QAM channels.

Implementing switched digital video as a method to preserve and expand digital bandwidth for high-definition and other services is one of several alternatives. You should conduct a similar cost characterization exercise for each of the other alternatives, including bandwidth expansion to 1 GHz, and full analog reclamation, which requires a $35 to $50 "dongle" for each TV in a home.

The fixed costs of SDV -- servers, QAMs, and all other vendor-supplied materials -- run in the range of $9 to $10 per home passed. Variable costs, like laser resegmentation, node splitting, and miscellaneous costs, can total $3 to $5 per home passed. Thus the total cost -- fixed and variable -- to implement SDV is in the range of $12 to $16 per home passed.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **www.cisco.com/go/offices.**