SIPV Platform Frees MSOs to Pursue DOCSIS Expansion without Upgrades

Accelerating Adoption of Adara Solution Underscores Game-Changing Impact on Broadband & TV Strategies

Table of Contents

Introduction

Growing numbers of cable MSOs are ensuring they'll have all the bandwidth they need for DOCSIS expansion now and for years to come by deploying a switched video platform that far outperforms much costlier approaches to increasing broadband access rates.

Pioneered by Adara Technologies, Switched IP Video (SIPV) is enabling MSOs large and small to implement DOCSIS bandwidth expansion with unprecedented flexibility, speed and performance while freeing them to add unlimited numbers of HD and, in the future, UHD channels, with no disruption to subscribers. Their successes are providing convincing evidence that by revamping and repurposing the SDV (switched digital video) technology that has long been a mainstay of cable operations, Adara has made it possible for operators to pursue their highest priority goals without shouldering the costs, complexities and other drawbacks of traditional SDV.

Or, as CableLabs CTO Ralph Brown put it recently, Adara has "reinvented SDV."

In a nutshell, SIPV exploits the dynamic power of video switching technology to minimize the amount of bandwidth needed for delivering linear cable channels, freeing up massive amounts of bandwidth for DOCSIS. At the same time, SIPV makes it possible for operators to continually expand their channel lineups no matter how many HD and UHD channels they might want to add over time without adding more video bandwidth.

All of this can be accomplished with SIPV at far lower costs with better results than node splitting, RF capacity expansion, overbuilding with FTTH or introducing new compression technology with new set-top boxes, which explains why ever more operators are turning to SIPV as their priority recourse. In so doing, they also are ensuring they will get better results from the other methods, if and when they ever need to be implemented.

The availability of the SIPV platform comes at a crucial time for cable operators. While they are in a powerful position as proprietors of the best broadband access network infrastructures in North America, they are hard pressed to keep up with demand for ever more broadband bandwidth when it comes to funding and executing evolution of the physical HFC plant. As a software-based mechanism for controlling use of existing network capacity, SIPV provides MSOs the tools they need to compete cost effectively at the pace of developments in the Internet era.

Thus, for example, an MSO with a limited amount of bandwidth available for DOCSIS owing to the volume of bandwidth used to deliver analog and digital TV services, doesn't have to invest in plant expansion or analog reclamation to make room for DOCSIS 3.1 if it deploys SIPV. Hundreds of MHz of RF spectrum can be immediately freed for DOCSIS 3.1 by using the advanced switching mechanisms Adara has developed to accommodate delivery of all the operator's digital TV content, thereby reducing the number of QAM channels used for TV to a fraction of current levels.

In the discussion that follows we first will look at some of the key market indicators that are shaping network evolution strategies in an environment that has already driven broadband bandwidth consumption far beyond anyone's expectations. Looking at these indicators, it's clear there's no end in sight.

We'll then explain in detail how SIPV provides the optimal framework for addressing immediate and ongoing transport capacity requirements arising from these trends. As shall be seen, no matter what stage of plant evolution an operator might be in, by making implementation of SIPV the next step in bandwidth expansion they will be better positioned than ever for whatever the future holds.

Readers can also click on this link to view a video overview of SIPV that will prove helpful in understanding the basics of the platform.

The Unrelenting March to Ever More Bandwidth Consumption

Snapshots of Current Access Speeds in North America

From the dawn of broadband services cable operators have outperformed competitors in meeting market demand for ever higher access rates with top bandwidth tiers now reaching 500 Mbps and even 1 Gbps in some markets. This has had a profound impact on consumer experience both within and beyond the cable realm as competitors try to keep up and, in some cases, leap ahead.

The latest FCC report on broadband fixed access trends found the median download speeds experienced by broadband subscribers (averaged across all analyzed speed tiers) jumped 22 percent in 2016 from the year before to 39 Mbps, representing close to a four-fold increase from 10 Mbps in March 2011, when the FCC began generating its broadband reports.[1] The FCC found the most popular tiers are now in the 100-300 Mbps range, depending on what's offered by any given ISP, compared to 12-30 Mbps in 2011.

In Canada, where there's nothing comparable to the bandwidth reports put out by the FCC, the Canadian Radio-television and Telecommunications Commission (CRTC) has issued a ruling calling for availability of broadband access everywhere at 50 Mbps downstream/10 Mbps upstream by 2021 while implementing mechanisms to generate up to $750 million in funding to support network expansion in rural areas. According to the latest quarterly broadband usage report issued by CDN (content delivery network) operator Akamai, average speeds at which all Canadian users accessed Internet content from its CDNs in Q3 2016 was 14.9 Mbps, up 13 percent year-to-year.[2]

Clearly, whatever the actual usage experiences are from one region to the next, these numbers suggest people in general are enjoying unprecedented levels of broadband access. But, based on these trend lines, it's equally clear operators have no reason to believe the limits of what people will demand in the future have been reached. Consequently, while many operators may not be feeling immediate pressure to raise their broadband rates, all operators are under pressure to determine how they will stay ahead of the trend line going forward.

Determining What Comes Next

Nielsen's Law and the Emergence of Gigabit Access

There are many indicators to consider in determining what the future bandwidth expansion curve might look like. One is the generic, informal benchmark for growth known as "Nielsen's Law," articulated in 1998 by Konrad Nielsen, a distinguished engineer at Sun Microsystems and now principal in the engineering consulting firm Nielsen Norman Group. His widely followed formulation that top-tier broadband bandwidth levels grow by 50% per year has proven to be an uncannily accurate predictor of his own broadband service options as well as bandwidth rates in general, starting from his first ISDN connection at 300 baud in 1983 and extending to 2016, when he upgraded his bandwidth to 240 Mbps.[3]

Extending the application of Nielsen's Law into the future, using Nielsen's 2016 top-tier 240 Mbps access rate as the starting point, sets the top-tier data rate target at about 1 Gbps in the 2019-2020 timeframe, at 10 Gbps around 2025 and at a hard-to-imagine 75 Gbps by 2030. Industry experience shows that the most popular mid-range broadband tiers hit any given billboard rate within a lag time of two to three years, which, when applied to the Nielsen Law projections, would mean 1 Gbps would be registering as a mainstream popular tier by 2022-2023.[4]

While no one can accept such numbers at face value for network planning purposes, there are plenty of reasons to give them credence based on current developments. One sign of where the market is headed is the quickening pace of 1 Gbps broadband service rollouts. By 2020, Deloitte predicts 600 million subscribers worldwide will be on networks offering a 1 Gbps service.[5]

Since Google Fiber made its debut in Kansas City in 2012, there has been an outpouring of 1 Gig launches in the U.S., including a few by Google, which recently said it was getting out of the network operations business, but many more by over-builders and incumbents. AT&T, for example, says it is offering 1 Gig service in 29 cities – a good share of which have more than one provider offering 1 Gig service – with another 38 projects in various stages of planning or construction.[6]

Local government- and utility-backed projects have been growing as well. By one count, 1 Gig fiber networks funded by tax payers can be found in 15 U.S. towns and cities.[7]

The OTT Video Factor

The most powerful driver to ever more broadband bandwidth consumption, of course, is video. Overall, global IP video traffic is on course to grow at a 26 percent CAGR over the next few years, reaching a level representing 82 percent of all consumer Internet traffic by 2020, according to predictions from the latest Visual Network Index compiled by Cisco Systems.[8]

Just how rapidly consumers are embracing the use of connected devices to view OTT live and on-demand long-form video is reflected in a number of studies. For example, based on metrics gathered from hundreds of millions of users worldwide, online video publisher Ooyala recently reported long-form video consumption on tablets and smartphones went from 25% of the time spent watching video on tablets and 16% on smartphones in Q4 2014 to 46% and 25%, respectively, in Q4 2015.[9] The share of long-form video viewing on PCs jumped over that same period from 19% to 32%.

The TV is also playing a big role in the growth of online video consumption, due to the proliferation of smart TVs and a wide range of IP streaming media players that have made OTT content available to TV viewers. In the U.S., at of the start of 2016, 45% of U.S. broadband households owned at least one smart TV (up from 34% a year earlier) and 36% owned at least one streaming media player (compared to 27% a year earlier) according to Parks Associates.[10]

By 2019, smart TV penetration in the U.S. will reach 57%, according to projections formulated by IHS.[11] Cisco's VNI study predicts Internet video-to-TV traffic globally will represent 26 percent of consumer video traffic by 2020, marking a 3.6-fold increase from 2015.

Working in tandem with video-capable device proliferation to drive OTT video consumption is the growing number of subscription options available to online viewers. Joining long-standing SVOD providers like Netflix, Amazon, Hulu, and Vudu is an increasing population of OTT skinny pay TV bundlers, including MSO's archrivals Dish with Sling TV and AT&T-owned DirecTV with DirecTV Now.

At the same time, many traditional broadcast and premium TV networks are taking action independent of MVPDs to keep pace with consumer demand.[12] Some, like HBO with HBO Now, NBC with distinct OTT options, and CBS with CBS All Access, are marketing standalone subscriptions to their content. Others are relying on advertising or hybrid ad/subscription models.

These OTT offerings are impacting bandwidth consumption in pay TV as well as broadband-only cable subscribers. According to Parks Associates, 52% of U.S. pay TV households now subscribe to at least one OTT service.[13] Overall, 64% of broadband households now subscribe to an OTT video service, Parks said.

Compounding the impact on bandwidth consumption is the fact that ever more devices are being used by household members to access video programming, often at the same time. Based on tracking of actual network usage in 2016, Sandvine, a supplier of broadband traffic-optimization equipment, said the average North American household had at least seven active devices in use daily, with video streaming accounting for 65% of usage.[14] By 2019, the average number of connected media-enabled devices per North American household will climb to ten, with the global average hitting five, according to IHS projections.[15]

The Upstream Video Impact

Bandwidth pressure from all these devices extends to upstream as well as downstream usage. With billions of smartphones, tablets, and PCs equipped with high-resolution cameras, video calling, the use of video and photos in social networking and uploads to cloud storage are consuming ever more bandwidth.

This is one reason the cable industry is pushing ahead with preparations for Full Duplex (FDX) DOCSIS 3.1, which will enable a symmetrical broadband service at speeds of up to 10 Gbps. CableLabs is preparing recommendations for best approaches to achieving FDX in anticipation that MSOs will begin introducing FDX services by 2019.[16] FDX relies on a combination of passive HFC and the self-interference cancellation and intelligent scheduling of DOCSIS 3.1 technology to enable use of the same spectrum for simultaneous downstream and upstream bit flows. The 10 Gbps goal points to expectations that operators will be using 1.2 GHz of RF spectrum to deliver broadband services in just a few years.

The Marketing Incentive to Offer Faster Broadband

From a marketing and sales standpoint, it's also important to note that, even if an operator sees no immediate need to raise the throughput on the DOCSIS network, there's an important benefit that comes with adding a faster top tier stemming from what is known in marketing circles as the "Compromise Effect." The concept rests on the principle, well established in retail pricing and academic studies, that people who might normally choose the lower priced of two options or go either way on a more or less 50/50 basis when the difference between them isn't clear, tend to choose the higher priced option if a third one is added at an even higher price.

As shown in Figure 1, this is the theory that inspired the three-price scheme based on octane levels that characterizes how gasoline is sold at gas stations just about everywhere. People tend not to see what the value is in choosing the highest octane option when presented two or three choices, but if there are three choices, they tend to go for the compromise middle option as a safe bet. Studies have found that adoption of this approach by the industry as a whole led to a 60% increase in the mid-tier octane option.

Figure 1

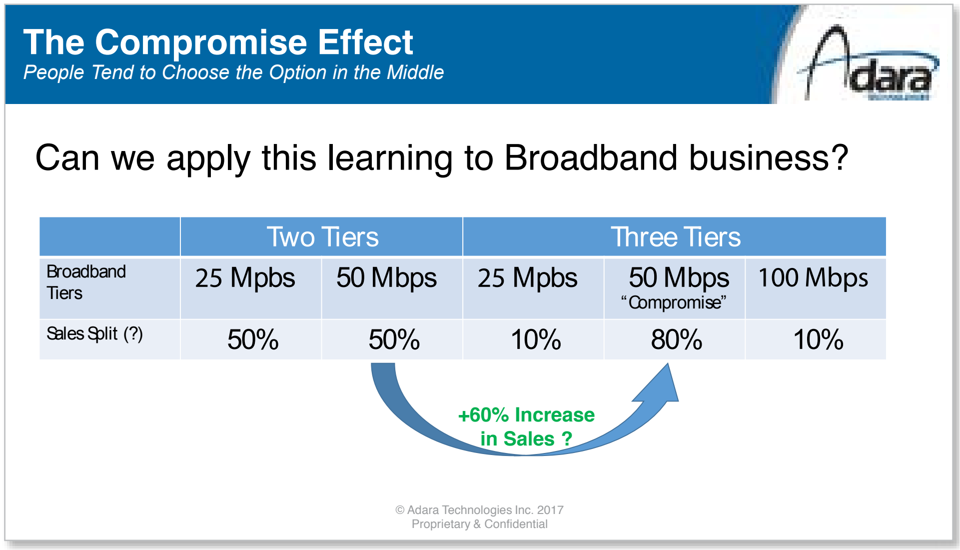

The lesson for cable operators, which many have already applied in broadband tiering strategies, is that adding a "billboard" very high-speed tier to the current options will generate more sales of what was the highest tier. People may not understand what the difference in their experience will be between, say, a 25 Mbps and a 50 Mbps option, leading them to choose 25 Mbps, but if a 100 Mbps option is added, 50 Mbps looks more like the way to go.

The Imminent Impact of 4K UHD on Bandwidth Requirements

Another big driver behind ongoing cable bandwidth expansion is the emergence of 4K Ultra HD. No matter what degree of confidence operators may have that current DOCSIS offerings will satisfy demand for some time to come, 4K UHD will force them to accommodate the need for more transport capacity on their systems if they want to avoid having to reduce the allocations to DOCSIS.

4K UHD has jumped to the front burner on the strength of two major developments: the industry's incorporation of HDR (high dynamic range) into 4K UHD as the real differentiator when it comes to delivering a compelling new viewing experience and the soaring penetration of 4K UHD-enabled smart TV sets, most of which now on offer in retail stores are equipped to support HDR.

Researcher IHS Markit, citing plummeting prices, recently predicted household penetration in the U.S. will reach 34% in 2019.[17] According to Futuresource Consulting, worldwide shipments of 4K UHD sets will account for 52% of the market by 2020, at which point the vast majority of television households in North America, Western Europe, Japan, South Korea and many other parts of the world will be watching TV on the big displays.[18]

There appears to be growing consensus within the TV industry that there's an ROI upside to offering 4K UHD services. A recent global survey conducted by SNL Kagan found that 64% of MVPDs and 73% of content producers among the nearly 500 respondents believe consumers will be willing to pay 10% to 30% more on their subscriptions for access to 4K UHD content.[19] Ninety-six percent of all respondents believe 4K UHD TV services will be widely adopted by 2020.

MSOs are feeling intensifying pressure from satellite and OTT competitors to make HDR-enhanced 4K UHD services available to their subscribers. In the OTT domain Netflix and Amazon have led the way so far with the addition of HDR-enhanced programming to 4K portfolios they've been building since 2014. Each of these providers said it was offering about 150 hours of HDR UHD content entering 2017.[20]

Other services like Vudu, Fandango Now, Sony Ultra, Hulu and UltraFlix4K are populating the Internet with ever more UHD options, some with HDR. And YouTube, which over several years has built a sizeable library of 4K content, has brought long-form HDR-enhanced UHD programming into the mix of channels offered on its Red subscription platform.[21]

So far, DirecTV has led the pay TV market in UHD with three channels currently devoted to the format, two on a 24/7 basis, the other for pay-per-view, and satellite capacity in place to support as many as 50 channels once the content becomes available. A formidable component of the DirecTV programming has been live sports coverage in 4K, which began with the Masters golf tournament in 2016, repeated with two-channel coverage in 2017, and now also includes occasional broadcasts of live MLB, NBA and Notre Dame football games. The DBS operator recently shifted from making UHD channels available at a high premium to other services to including UHD as part of its 145-channel, $50-per-month "Select" plan.

Contemplating the Future Implications of Virtual Reality

While it's too early in the evolution of virtual reality to predict its impact on cable operators' service offerings, the emergence of this bandwidth-hungry format underscores the likelihood that the upward trajectory in demand for throughput will continue for years to come. Improvements in head gear, better production technology and growing volumes of content suggest the ramp up to a viable VR market might not take as long as it once seemed.

Goldman Sachs Group recently predicted an $80-billion global market for VR and AR (augmented reality) will emerge by 2025 with $45 billion generated from hardware sales and the remainder from various applications categories, including $7.4 billion from live event and episodic entertainment in addition to $11.6 billion from games.[22] While games have been the dominant content category so far, there's now significant movement on other fronts as well.

Entities now engaged in regular VR content production include Discovery Communications, The New York Times, CNN, the BBC, Netflix, DirecTV and other major media outlets. A recent survey of Hollywood content creators jointly sponsored by the Consumer Technology Association and NATPE (National Association of Television Program Executives) found that most respondents believe VR represents a game-changing method of storytelling.[23] Sports and concerts were also cited as promising areas of development.

Significant progress has been made on modes of transmitting VR content over networks that reduce the bitrate to include just the frames viewed by each eye at each instant in time. The effectiveness of this approach has been demonstrated in live broadcasts of professional auto racing, golf, basketball, soccer and hockey, along with other types of events. But the need to transmit 4K-size displays for each eye at a frame rate of 90 Hz using HEVC encoding means it will take on the order of a 100-150 Mbps throughput to support fully immersive VR, according to CableLabs.[24]

Finding a Cost-Effective Future-Proof Path To DOCSIS Expansion

Clearly, no matter what level of broadband service an MSO might be offering today, there can be no let-up in planning for best approaches to meeting the need for more DOCSIS bandwidth. Compounding the challenge is the fact that operators will need to accommodate ever more TV carrying capacity as well.

The Limitations of Traditional Methods

But operators need a more cost-effective, scalable way to deal with these challenges than they've had in the past. Most have already performed analog reclamation, which has more often than not served to meet goals on the TV side of the house resulting in the filling of the freed RF spectrum with more HD broadcast channels and unicast channels for surging levels of VOD usage. In instances where analog reclamation has yet to be executed, operators can expect the same types of outcomes, resulting in little, if any extra bandwidth for DOCSIS.

Installing new amplifiers and pushing fiber deeper to increase RF spectrum capacity on the coaxial plant and splitting nodes to lower the number of people competing for available bandwidth have long been fundamental steps in the push to make more bandwidth available for DOCSIS. But the gains achieved through these measures suffice for only so long before the whole process has to be repeated, at great cost and disruption to operations.

Operators can also free up bandwidth through deployment of set-top-boxes (STBs) supporting higher levels of compression. But typically any such steps are tied to upselling subscribers to a new tier of TV service, resulting in an incremental per-household reduction in bandwidth consumption that has a negligible impact on broadband spectrum availability system wide. A system-wide change-out to the new STBs can free up some spectrum for broadband, but this incurs massive CPE and installation costs as well as inconvenience to subscribers, which is why most MSOs don't pursue this course.

Another recourse in the search for more broadband throughput, DOCSIS 3.1, has provided the industry a far more efficient way to use available RF capacity compared to previous generations of DOCSIS through techniques such as the application of Orthogonal Frequency Division Multiplexing (OFDM) across channels consuming 192 MHz each and raising QAM levels to as high as 4096. DOCSIS 3.1 can achieve what can be up to a 60+% gain in bits per Hertz versus DOCSIS 3.0 when QAM 4096 is used or something in the range of 30% with lower levels of QAM.

But, DOCSIS 3.1, too, is a costly undertaking, and, in any event, whenever it is deployed it will deliver far more throughput if the operator can dedicate more spectrum to its use than might be available now. One way of getting that additional spectrum is to spend even more money on previously discussed options. Another is to free up more spectrum without resorting to any of those methods, which is where the Adara Switched IP Video (SIPV) solution comes into play.

The Dramatic Impact of Expansion Strategies Based on Adara's SIPV

Because SIPV provides far more broadband bandwidth to work with whether or not any of these other steps are taken and at far lower costs than would be incurred with these other measures, it becomes the obvious next step to DOCSIS expansion. That it can do this while providing support for unlimited expansion of video channels greatly strengthens the case, and it terminates the ongoing battle between the video and broadband sides of the business over how available bandwidth is divided between them to meet competitive and contractual obligations.

Figure 2 offers a vivid illustration of the impact SIPV has on bandwidth allocation in any existing plant situation. Without SIPV (left image), most of RF spectrum capacity is devoted to carrying TV signals, which can include any combination of linear analog and digital HD and SD channels as well as capacity for VOD, leaving relatively little spectrum for DOCSIS. With SIPV in place (right image), just a small portion of the RF capacity is used to carry all the digital video, leaving most of the bandwidth for allocation to DOCSIS either all at once or incrementally over time.

No wonder, then, that SIPV has become the logical next step in network evolution for ever more cable companies, including large Tier 2 as well as Tier 3 and 4 operators, no matter how far they've gone with other techniques. Indeed, most of Adara's customers have already employed at least one of the other approaches to expanding bandwidth for DOCSIS. They understand that not only does SIPV offer a much better way to address immediate needs; it provides a foundation for getting much more out of more costly plant upgrades if and when they become necessary.

The ability to employ on-the-fly switching technology to add efficiencies to bandwidth utilization has long been a mainstay in cable operations as a fundamental component of how DOCSIS works on the broadband side and how SDV functions, on the TV side. Video switching is also essential to how DSL works as a conduit for delivering TV services via IPTV.

But in cable, apart from the limited use of SDV for enabling delivery of more long-tail linear content to subscribers, the entire subscriber base in any given cable system is treated as a single service group (SG) receiving all linear channels at once in broadcast mode, even though only 20% to 30% of those channels are watched at the highest levels of peak usage. See Figure 3 for a summary of the differences and similarities between SDV and SIPV.

Leveraging SIPV to Maximize DOCSIS Bandwidth

As illustrated in Figure 4, SIPV eliminates this inefficient use of bandwidth by applying switching technology to the entire channel lineup or any portion thereof, ensuring that only those channels being watched by one or more users within a SG are multicast over the network. The platform places channels into available slots on the QAMs that are dedicated to SIPV as those channels are accessed by viewers, freeing operators from bandwidth constraints on the number of video channels they can support while making hundreds of MHz of bandwidth available for DOCSIS expansion on existing plant.

There is no change in the way TV signals are distributed to set-tops, which is to say, the packets of digital video are encapsulated for distribution over MPEG-2 TS to legacy STBs, requiring no change in CPE or other aspects of operations. Like SDV, SIPV maintains the quality of the incoming program stream, since it does not compress or modulate it differently.

Going Far Beyond SDV

In creating SIPV, Adara has leveraged its long experience as the supplier and operator of cloud-based technology to create a solution that utilizes existing SDV platform components as the foundation for switching unlimited video sources in a small part of the spectrum without onerous plant investments or STB change outs. Thus, the platform relies on existing SDV client software to support tuning to a channel that already is in distribution over one of the QAM slots used by SIPV or to trigger distribution of an unwatched channel over an unused QAM slot. Because SDV client software is embedded in the vast majority of STBs in the field, most terminals already have all that's needed to support SIPV.

At the same time, as shown in Figure 5 and discussed at length below, there is much that differentiates SIPV from SDV. Adara has employed its cloud-hosted service model and

Figure 2

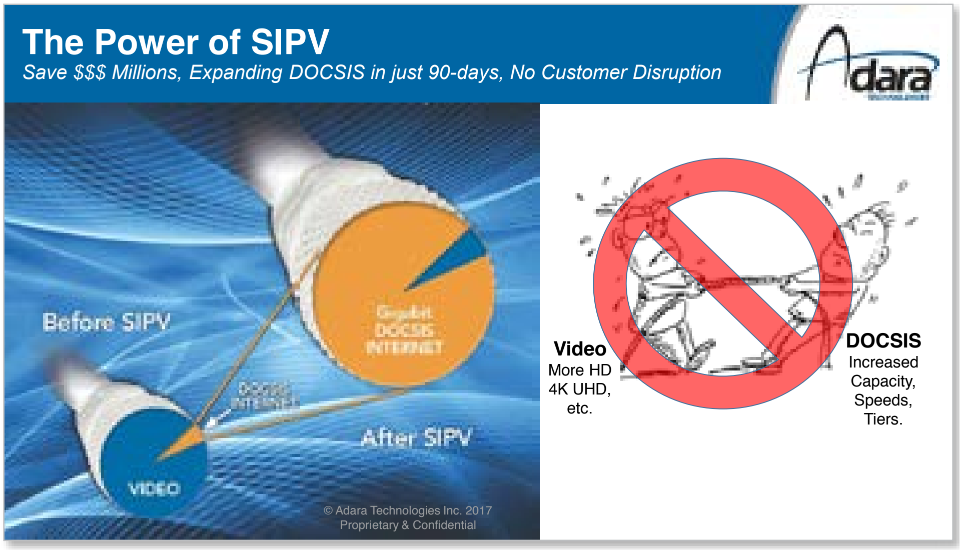

Figure 3

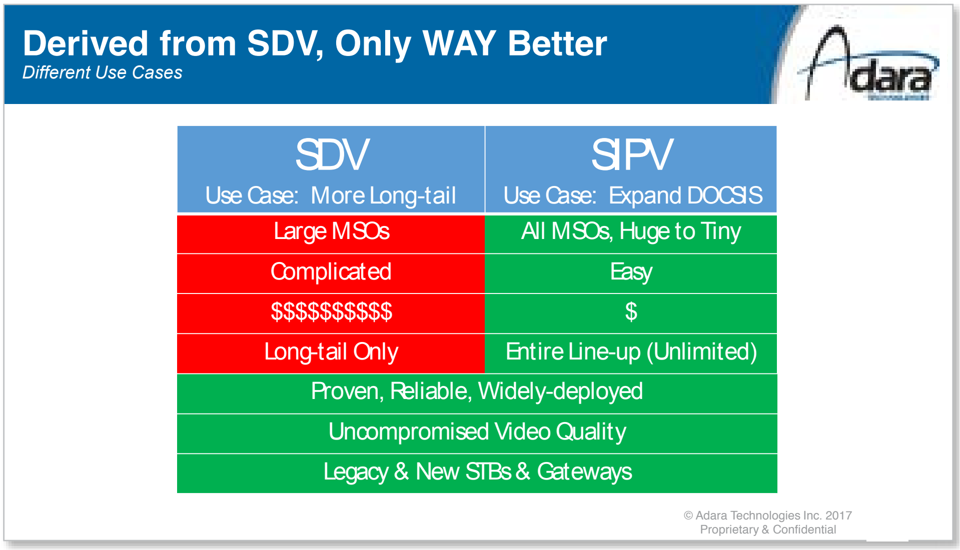

SDV Use Case: More Long-tail	SIPV Use Case: Expand DOCSIS
Large MSOs	All MSOs, Huge to Tiny
Complicated	Easy
$$$$$$$$$$	$
Long-tail Only	Entire Line-up (Unlimited)
Proven, Reliable, Widely-deployed	
Uncompromised Video Quality	
Legacy & New STBs & Gateways	

Figure 4

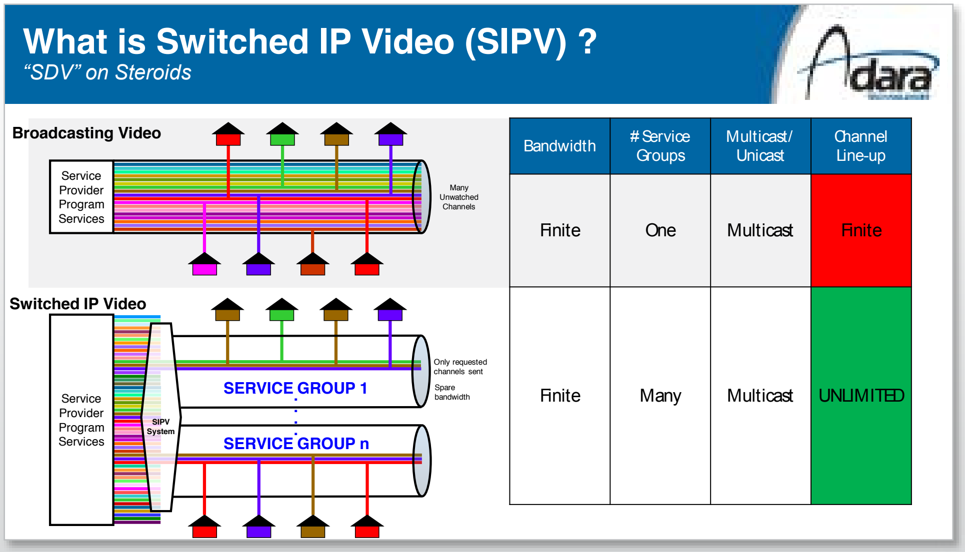

	Bandwidth	# Service Groups	Multicast/ Unicast	Channel Line-up
Broadcasting Video	Finite	One	Multicast	Finite
Switched IP Video	Finite	Many	Multicast	UNLIMITED

Figure 5

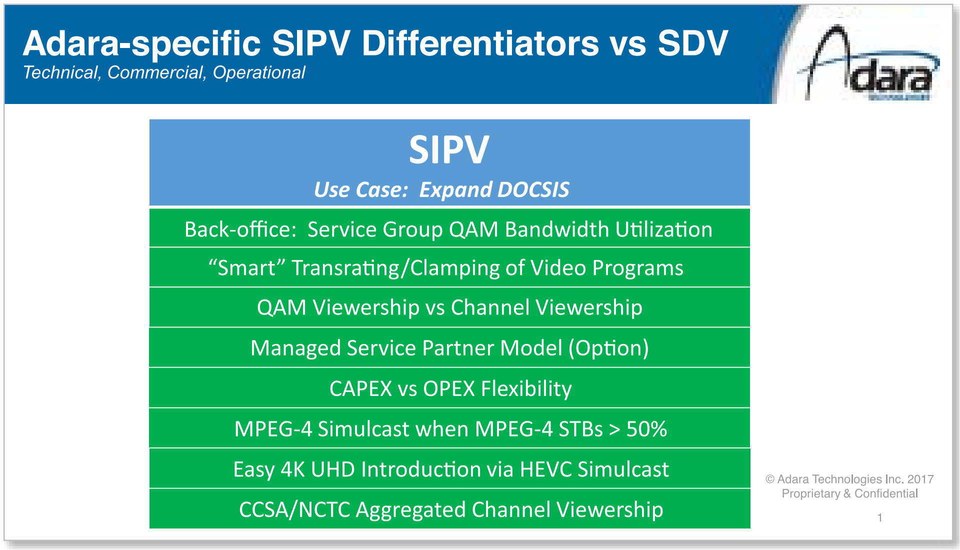

Figure 6

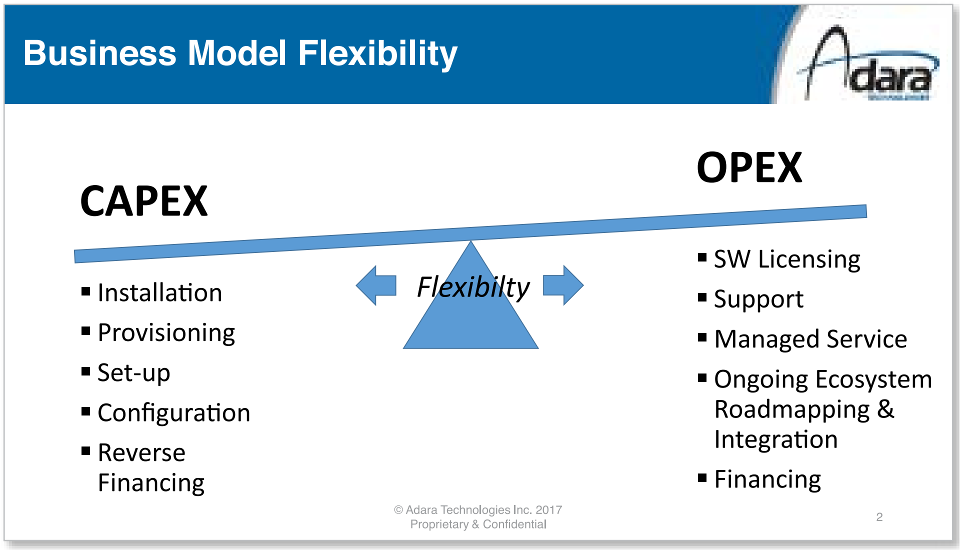

Figure 7

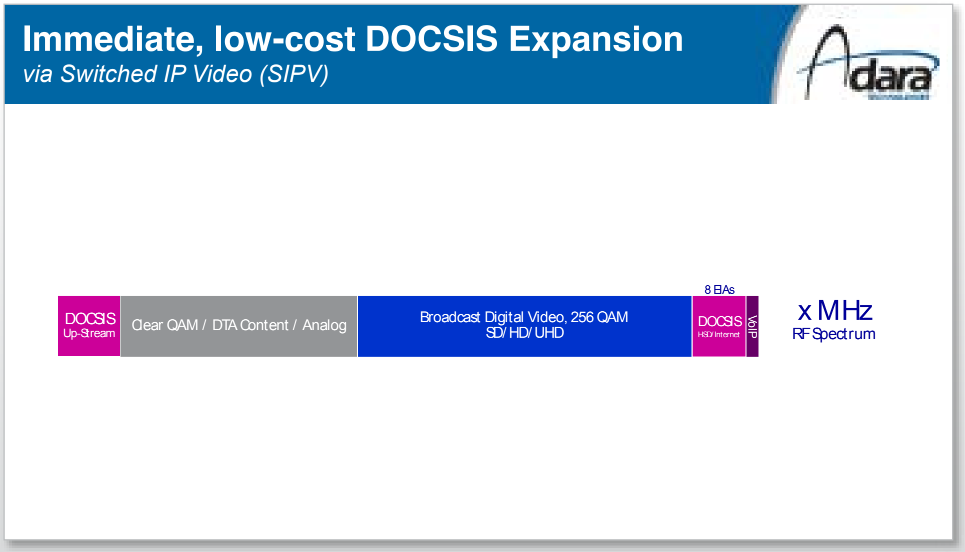

other innovations with SIPV to avoid traditional SDV encumbrances, including big capital costs and the OpEx costs associated with fielding teams of operations and engineering personnel. Moreover, SIPV brings into play unique software utilities that are essential to adding ever more HD and UHD channels without impacting the DOCSIS bandwidth allocation.

If the MSO has an analog channel lineup that it wants to retain with SIPV deployment, the amount of RF spectrum dedicated to those channels will remain the same. Some operators using SIPV have retained their analog channels, having found the platform frees up all the DOCSIS bandwidth they need without reclaiming the analog spectrum.

Like SDV, SIPV divides the video subscriber population into SGs employing the headend mechanisms that are used to combine multiple node service areas into SGs for DOCSIS and VOD. At the outset of SIPV implementation the operator and Adara determine how many QAMs will be assigned to modulating 6 MHz EIA (Electronic Industry Alliance) channels within each SG.

Typically, the initial SG sizes in terms of the number of node serving areas within each will be based on multiples of four, in line with how DOCSIS SGs are configured.

A Flexible Approach to Setting the Pace of DOCSIS Expansion

One of the great benefits of SIPV is the flexibility it affords operators in their approaches to accommodating DOCSIS expansion. As illustrated by Figure 6, this flexibility extends to enabling each operator to utilize SIPV in whatever ways comport with its approach to budgeting for costs, whether the preference is to rely on CapEx funding or to accommodate costs as part of ongoing OpEx outlays.

In Figures 7-9 we provide a view of how an operator might progressively use bandwidth freed up for DOCSIS by SIPV to accommodate various stages of broadband offerings over time, since, in many cases an operator may see no need to use the full RF capacity to achieve the full potential of its liberated broadband spectrum at the outset.

Figure 7 sets the stage as an illustration of what the QAM allocations for all services might look like before SIPV is deployed. In this typical scenario most of the RF spectrum on a plant supporting anywhere from 550-1,000 MHz of total spectrum per node is consumed by analog and/or clear-QAM/DTA content and broadcast digital video spectrum, leaving just eight EIAs or 48 MHz of spectrum to support downstream broadband throughput along with the usual spectrum allocations for upstream DOCSIS and digital voice.

In this and subsequent illustrations we're showing a high ongoing allocation to analog or one-way DTA channels. Larger Tier 3 and Tier 2 operators may have significantly reduced but perhaps not totally eliminated analog/DTA channels in order to continue providing OTA and other highly popular channels to subscribers having legacy one-way STBs or clear-QAM tuners or who have a TV or two with no STBs at all.

As shown in Figure 8, with initial deployment of SIPV anywhere from 8 to 24 QAM EIAs have been dedicated to carrying SIPV switched TV channels. This scenario assumes the operator retains the analog/DTA allocation with continued support of DOCSIS on the legacy broadband spectrum, leaving the vast amount of freed spectrum for DOCSIS expansion.

The number of EIA QAMs dedicated to SIPV depends on how the operator chooses to balance a number of inter-related considerations. These include the size of the SG served by SIPV, how many TV channels the operator wants to allocate to the SIPV EIA QAMs at the outset, how the channel mix breaks down with regard to numbers of SD, HD and 4K UHD channels, the relative popularity of the chosen channels and the proportionate allocation of channels to MPEG-2, MPEG-4 and HEVC.

The weighting of these different considerations can produce a wide range of outcomes, as evidenced in current SIPV deployments. For example, there are instances where operators have started out with allocation of 24 EIA QAMs to SIPV across SGs with, say, 1,000 subscribers each to deployments where 8 EIA QAMs are dedicated to serving SGs with far fewer subscribers.

With this level of flexibility in mind, an operator might determine that 8 SIPV QAMs are enough for now, even though planning projections suggest that, for the SG as defined in the initial set-up, the operator may eventually need as many as 24. In fact, depending on total spectrum available from each node and the number of EIAs dedicated to SIPV, the additional EIAs available for DOCSIS expansion beyond the initial eight could total 50 or more.

At the stage of the SIPV deployment illustrated in Figure 7 the cable system has been prepared for inevitable DOCSIS expansion to come without requiring any change in status quo offerings. Of course, the operator is free to use as much of the liberated RF spectrum to expand broadband tiers immediately to whatever levels market conditions merit. Whenever expansion occurs, customers have been able to go from the legacy scenario to being fully operational with SIPV within 90 days, with no disruption to subscribers.

Figures 9-10 show progressive allocations

Figure 8

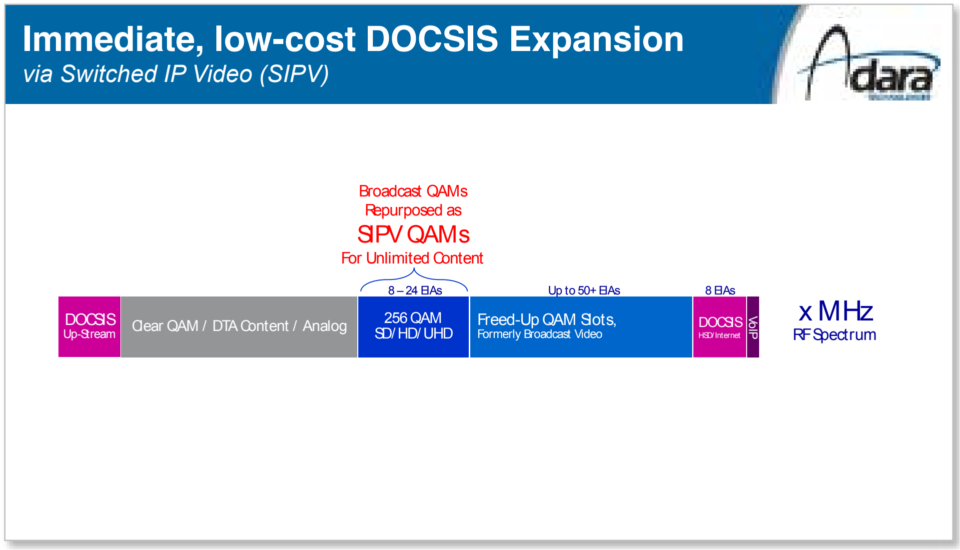

Figure 9

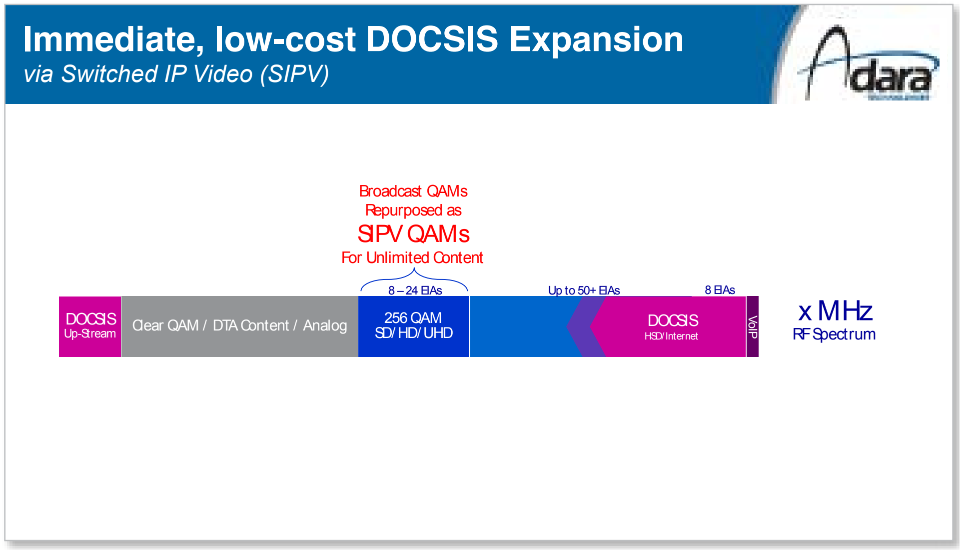

Figure 10

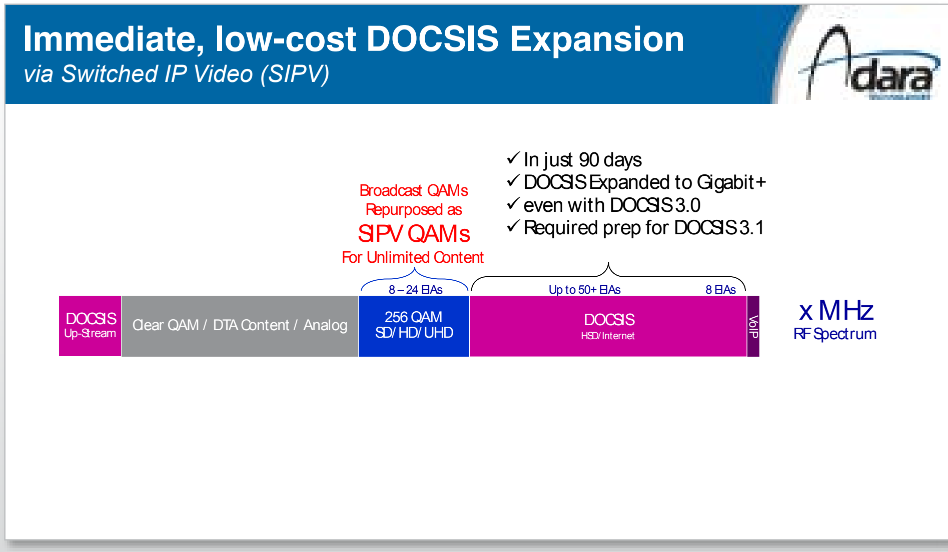

of spectrum freed up by SIPV to activation for DOCSIS as needs arise, culminating in dedication of the entire freed up capacity to DOCSIS. Freeing up and activating incremental spectrum for DOCSIS following initial implementation of SIPV takes just a matter of days.

Looking at Figure 10, if this is a 750 MHz plant and the EIA allocation to SIPV already totals 24, and if there are 24 analog/DTA channels consuming 144 MHz, the total additional spectrum available for DOCSIS on top of the legacy allocations for downstream and upstream DOCSIS and voice would be in the neighborhood of 366 MHz, or about 60 EIAs.

If the allocations were 24 SIPV EIAs and 60 analog/DTA channels, the gain in DOCSIS spectrum would add up to about 150 MHz. Thus, even in this scenario with a large analog channel total left intact, new DOCSIS spectrum would be equivalent to more than three times the amount of legacy DOCSIS spectrum.

The Implications for DOCSIS 3.1

Of course, virtually any amount of freed up spectrum can be used to support migration to DOCSIS 3.1 in both the downstream and upstream paths, insofar as downstream OFDM channel width can be as small as 24 MHz, with a minimum of two channels required no matter what their size, and minimum width for upstream channels operating in the OFDMA (Orthogonal Frequency Division Multiple Access) mode specified for DOCSIS 3.1 upstream is just 6.4 MHz, with, again, a minimum of two channels required.

In the examples used here, with DOCSIS 3.1 throughput over 4096 QAM totaling 8.3 Mbps per MHz of spectrum, total downstream throughput available over the 366 MHz of freed up DOCSIS bandwidth in the first instance comes to about 3 Gbps across the SG, leaving the legacy DOCSIS bandwidth to serve existing modems. In the second instance, the 150 MHz of freed up bandwidth would be sufficient to support 1.2 Gbps of DOCSIS 3.1 throughput. In more likely instances where QAM levels vary at different points in the spectrum, the throughput in these two instances might be in the neighborhood of 2.4 Gbps and 1 Gbps, respectively.

Not all this freed up spectrum need be allocated to downstream, of course. Operators have the option to allocate any portion to upstream that falls within the ranges allocated for DOCSIS 3.1 upstream transmissions, as long as the stipulation that two channels equaling a minimum of 6.4 MHz each is met.

Unlike the DOCSIS 3.1 downstream channels, which cannot share spectrum devoted to DOCSIS 3.0 or earlier generations, upstream channels can run on the same spectrum as the earlier versions. DOCSIS 3.1 can be implemented within the legacy 5-42 MHz sub-split, the 5-85 MHz mid-split or the maximum allocation out to the upper band edge of 204 MHz.

As these examples show, no matter what the spectrum allocations to legacy services might be, operators contemplating migration to DOCSIS 3.1 will be far better served by deploying SIPV first, ensuring they have the maximum bandwidth available for OFDM and OFDMA channels without any changes in the physical plant.

Building TV Service Expansion Strategies Based on SIPV

On the TV side, the volume of individual HD, SD and, eventually, UHD channels that can be supported by a given number of EIA frequencies on the SIPV platform is a function of the size of the SG and the level of peak usage within that SG. Of course, there are no "nailed-up" channel assignments when it comes to how SIPV QAMs are used, which means assessment of capacity utilization is done on an ongoing, aggregate basis. SIPV is about using switching technology to minimize the amount of bandwidth necessary to support the entire linear digital channel lineup within each SG, not merely to make room for seldom-watched long-tail channels.

The Distinctive Characteristics and Benefits of SIPV in the TV Domain

The Inclusion of All, Not Just Long-Tail Linear TV Content

This requires a different perspective on video switching than operators have used in applying SDV. SIPV maximizes total bandwidth efficiency for all channels by dynamically assigning channels to slots within the SIPV EIAs. In other words, whereas SDV is about selecting least-watched channels for placement into the pool of switched channels, SIPV is about dynamically assigning all channels whenever they're viewed to whichever EIAs have room for adding more channels at a given moment in time.

This isn't a random selection process. The SIPV software is continuously assessing available space and sequencing new channel assignments to optimum EIA slots based on where the fit is best for ensuring the maximum amount of available bandwidth within a given EIA is used with the new placement or will be used with the right mix of channels as the unused capacity is filled. Here the emphasis is on continually assessing how any given mix of SD and HD channels will lead to optimal capacity utilization.

Fast Channel Tuning

Another important thing to know about the SIPV experience is that it does not add any delay to the normal channel-tuning speed experienced with digital video insofar as adding or removing a channel from the multicast stream consumes just a few milliseconds. In fact, in some instances use of SIPV can actually lessen the time it takes for a STB to tune a new channel.

Of course, the basic tuning time, typically anywhere from 2 to 3 seconds, is a function of how long it takes the STB tuner to go from one frequency to another which varies according to the range of frequencies that must be traversed in the tuning process. Whereas in digital broadcast a given channel change may entail traversing hundreds of MHz of RF frequency, with SIPV the widest separation from the frequency where the current viewing is focused to the frequency where the chosen channel resides is just 24 EIAs, or 144 MHz, with average separation much less. This can reduce tuning time between any two channels by as much as half a second or more compared to the time consumed with tuning those channels in legacy broadcast mode.

Eliminating the Guesswork in Preparing for Resizing Service Groups

SIPV affords operators maximum flexibility to set up and adjust SG sizes to optimally match conditions specific to each SG. This means there's no requirement for making SG sizes uniform either within a given cable system or across multiple systems. Because SIPV manages QAM usage within each SG independently of how TV channels are placed in other SGs, SIPV removes the guesswork in resizing SGs individually in response to usage patterns within each group.

Another benefit operators gain from use of SIPV is the elimination of uncertainties from their efforts to determine when they will need to reduce the size of any given SG. As discussed at greater length below, the ability to anticipate and quickly act on the need to accommodate higher usage rates or more video channels by splitting SGs is enabled through the usage monitoring mechanisms of the automated Adara software utilities.

There are other ways in which flexibility in setting the sizes of SGs has great implications for maximizing the benefits of SIPV on a per-SG basis. For example, an operator might want to vary how many EIA channels are dedicated to SIPV within each SG as a function of current total RF bandwidth on that segment of the cable network. This can ensure that the same DOCSIS throughput is available across all SGs regardless of the spectrum differences often present across a cable plant.

Thus, in areas where capacity is at 860 MHz, it may become necessary to allocate 24 EIAs to SIPV with larger SGs while simultaneously allocating 16 EIAs to SIPV with smaller SGs in areas where 750 MHz plant is in operation. This means, with future video channel expansion, it may become necessary to reduce SG sizes in the 750 MHz areas but not in the 860 MHz plant. This built-in automated control flexibility ensures that the operator's broadband and video offerings can be uniform across non-uniform plant with a minimum of SG splits.

Or maybe there are great variations in terms of the number of subscribers served within each SG. In these instances, the operator might dedicate far fewer SIPV EIAs in low population areas than is necessary in more densely populated areas, leaving open the option to reconfigure those low-population SGs as needed with growth in population or gains in the subscriber penetration levels over time.

Also, even if plant upgrades or node splits become desirable at some point in the distant future, these investments can be minimized by virtue of the fact that they can be targeted to a very specific area or just one or a few SGs. All these possibilities highlight how SIPV, leveraging advanced analytics and independent management of QAMs within each SG, serves as a powerful bandwidth expansion tool distinct from traditional SDV.

Mounting and Executing SIPV

To illustrate how SIPV is implemented and SGs are optimized to enable support for an unlimited number of SD, HD and, eventually, UHD channels it's useful to reference a hypothetical scenario that is typical of the mid-size MSOs now using SIPV. Here we'll assume the MSO has allocated 16 EIAs to SIPV to support the current channel lineup in each of 20 SGs, which together comprise the entirety of the MSO's household base.

This will require assignment of 320 universal edge QAMs to SIPV at the outset. Typically, Adara supplies the U-EQAMs used with SIPV as part of its solution, although operators who are already using U-EQAMs in conjunction with CCAP sometimes choose to rely on their existing supply base. By definition any vendor's U-EQAM that meets industry specifications is compatible with SIPV.

The 320 U-EQAMs Adara would provide for this hypothetical MSO are compacted into three to four rack units consuming no more than 5-7 inches of rack space. These are installed at the headend or hub with no impact on the outside plant.

Another key headend component included in the Adara solution is the SIPV Universal Session and Resource Manager (USRM), which determines which SIPV QAM frequencies are assigned to channels in the first instance when a viewer in a given SG tunes to that channel. The USRM then automatically directs any other viewers of the same program within the SG to join the same multicast and "releases" the multicast from the assigned EIA frequency when there is nobody watching or recording the channel in that same SG.

The headend components of the SIPV platform also include Bulk Encryptors, which allow each linear channel to be encrypted independently of which U-EQAM is engaged whenever the channel is chosen for viewing within the SG.

With the ability to consolidate additional broadcast QAMs into SIPV EIAs there's plenty of capacity available to accommodate ongoing expansion of the channel lineup over time within the SGs as originally configured. But, as mentioned earlier, if that additional capacity eventually proves insufficient within one or more SGs, the Adara platform makes it easy for operators to determine when and how to reconfigure the impacted SGs.

Steps to Achieving Unlimited TV Channel Expansion

The role Adara's automated analytics utilities play in alerting the operator of a need to reduce SG sizes or add SIPV QAMs is illustrated by the screen capture from an operator's SIPV dashboard depicted in Figure 11. The graphic reflects Adara's continuous tracking of video consumption against available SIPV bandwidth within each SG.

Upon seeing evidence that the aggregate early-alert usage threshold set by the operator, in this case 80%, has been or may soon be breached, the operator and Adara begin cooperating on the processes they will undertake in tandem to reconfigure any SGs where there's a need to lower pressure on SIPV capacity.

The first process step in this task is communication between the operator and Adara on the need for and timing of a SG split, which typically cuts the SG size by half, resulting in a new SG requiring its own set of U-EQAMs. Once a date for the splitting is agreed upon, Adara creates a ticket for the job that allows the operator and Adara to keep each other appraised of progress over the course of the task.

At this point Adara orders and delivers the U-EQAM for the new SG in order to give the new SG the same quotient of EIAs to work with that has been employed by the original SG in the SIPV switching process. At the same time, the operator engages in pre-split preparations, starting with prewiring that will enable the unbundling and reconfiguration on the appointed date following arrival of the new U-EQAMs. The operator communicates to Adara with confirmations the U-EQAMs have been delivered, estimations of when they'll be racked and wired up and confirmations that those tasks are completed so that the provisioning of U-EQAMs for operation in the reconfigured SGs can be implemented on schedule.

In most cases operators prefer that Adara provides SIPV as a managed service, alleviating the operator of dealing with the complications

Figure 11

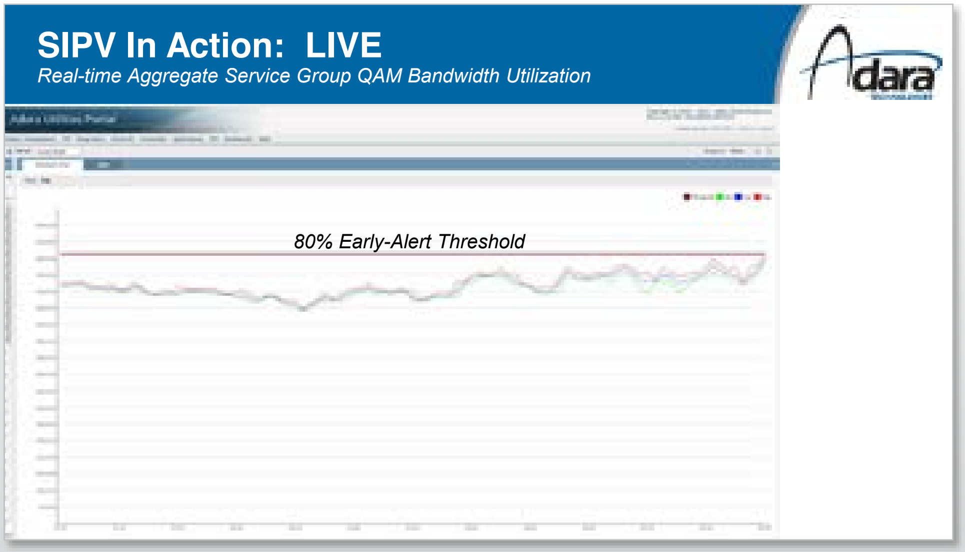

of the provisioning process. Once the new set-up is ready for activation, the newly configured SG can be put into operation within a ten-minute window during overnight off-hours, causing minimal disruption to the few viewers that might be watching TV at the time. The entire process, from receipt of the U-EQAMs to activation on the new SG topology, can take as little as a few days. Or more time can be allocated to preparations by the operator, usually lasting no more than a few weeks.

Implementing 4K UHD

As is clear from all the foregoing, SIPV provides operators with an extremely cost-effective way to introduce 4K UHD services. This solves a major conundrum for operators, who, as discussed earlier, are under increasing pressure to introduce 4K. The widely perceived problem is that this is too hard to do on the legacy TV side of the network and so will require use of broadband spectrum to set up an IPTV or OTT mode of delivering the new services. But this would not only mark a sharp departure from standing operating procedures; it would also result in increased operational complexity and major new investments in a new video control plane ecosystem.

SIPV provides the solution that keeps 4K in the mainstream of pay TV operations, requiring only the availability of 4K content that can easily be added or simulcast and the associated 4K-capable STBs that can be target-deployed only to paying subscribers of the 4K content. Additional bandwidth requirements are no longer a concern to operators as SIPV's SG bandwidth analytics continue to monitor utilization, and, more importantly, incremental loading of the DOCSIS network can be avoided.

While operators delivering 4K UHD will likely use HEVC or another next-generation compression technology supported by new STBs, the bandwidth efficiencies intrinsic to SIPV leave open the option to leverage MPEG-4 at least initially when there are a limited number of channel options and utilization is fairly limited.

In any event, with SIPV there's no need for a system-wide STB change-out to get started with 4K UHD. The platform recognizes whether the STB of a particular customer tuning to a 4K UHD channel is equipped with a 4K-capable STB and, if it is, instantly tunes the STB to the appropriate EIA frequency. Therefore an operator can immediately stand up a UHD 4K simulcast of, say, 20 channels and target deploy 4K-capable STBs to customers willing to pay for that new programming package.

With SIPV 4K UHD channel usage just becomes another contributor to aggregate QAM usage within each SG. When the maximum usage threshold is reached, the operator can begin the SG splitting process.

Leveraging Codec-Awareness to Maximize Efficiency in a Mixed STB Environment

Much as SIPV makes it possible to accommodate introduction of 4K UHD without having to allocate additional U-EQAMs, the platform can be used to mix channels targeted to MPEG-2 and MPEG-4/2 STBs, and MPEG-4/2/HEVC STBs when they come into play. Thus, experience in the field has shown that there's more bandwidth efficiency to be gained in a mixed MPEG-2/MPEG-4 STB environment by simulcasting the channel lineup allocated to SIPV in MPEG-2 and MPEG-4 when MPEG-4 STB penetration reaches 50% or greater within a given SG rather than delivering all programming solely in MPEG-2 until such time as the penetration of MPEG-4 capable STBs gradually (or via a costly and disruptive STB forklift) reaches 100%.

In such cases there's typically at least a 70% or greater chance that the initial tuning to a given channel within the SG will be performed by an MPEG-4 STB and that subsequent tunings to that channel will be by MPEG-4 STBs. This is especially true when one considers that, typically, the newer MPEG-4 capable STBs have been deployed on the most-watched TVs in the subscriber home, relegating the older MPEG-2 STBs to less-watched TVs. As the chances are reduced that the next STB tuning to that channel or any other channel in the SIPV lineup will require activation of the MPEG-2 version, the efficiencies of the simulcast approach increase reaching 30% or more in comparison to delivering all programming in MPEG-2 only.

The Role of Adara Analytics in Maximizing SIPV Benefits

Beyond tracking usage for SG management purposes, the analytics capabilities of the Adara SIPV platform provides operations, marketing and front-office personnel invaluable information they can use not only to understand what's happening in real time on their cable networks but also to look at usage trends as they plan for future needs.

Figure 12 offers a screenshot taken from the cloud-supported dashboard which an MSO uses to monitor what is happening across all the dynamically assigned channels currently being viewed on the SIPV QAMs within each SG. As mentioned, the USRM manages each SG independently, so the dashboard views of what's going on with the dynamic assignment of viewed channels to U-EQAMs vary greatly from one SG to another.

In this instance, the MSO has activated 24 EIAs in the SG reflected here, which consists of six bundled node serving areas representing a total of 652 STBs. The aggregate number of linear channels managed by this operator's instantiation of SIPV includes 350 SD and 150 HD channels, which means that about 660 MHz worth of spectrum that would be required in traditional operations has been reduced to utilization of only 144 MHz by SIPV.

Each column represents the activity underway at this moment in time on each of the EIAs. The amount of bandwidth consumed by each currently viewed channel is represented by the size of the vertical block named for the channel, with the maximum block size

Figure 12

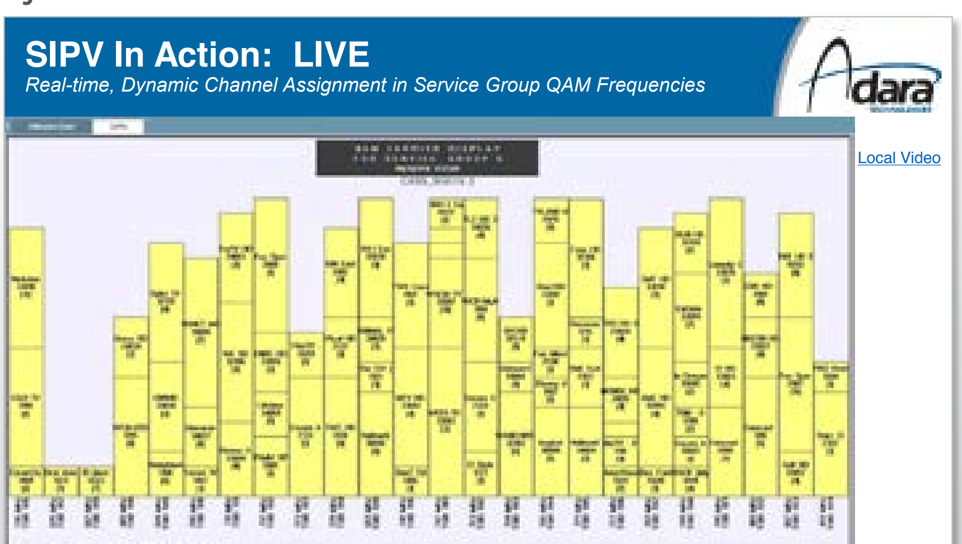

assigned to HD signals using MPEG-2 compression and the smallest size assigned to SD signals using MPEG-4. Within each block the number in parentheses accounts for how many STBs are tuned to the channel.

When the viewership on a given channel drops to zero, the system terminates the assignment of that channel to that EIA and the graphic shows the previously occupied space as white space available for use by another TV channel. As mentioned earlier, available spaces are filled as new channels are tuned to within the SG based on the optimal fit for the bandwidth consumed by that channel.

Figure 13 provides another SIPV dashboard view of actionable information generated by the Adara analytics system. Here the MSO has a comprehensive view of aggregate viewership for all channels for the current month, noting that only the 28 most-watched channels are shown on the screen and that scrolling will show the balance of 350 channels. The operator can export this chart and/or the raw viewership data to produce whatever additional charts are desired.

Many bases are covered by the information promulgated here. For instance, in the column immediately to the left of the bar chart the set of three numbers represents the total time (hours, minutes, seconds) that channel was watched by all STBs in aggregate for that particular month.

Such data provides operators important metrics with far greater detail than they can obtain from Nielsen ratings. For example, the data generated by the SIPV solution is specific to all the operator's subscribers and all channels. They can gain greater surety that they are offering people what they want to see, assess where interests are pointing when it comes to adding new channels and ensure they are getting fair value for what they're paying rights holders.

In instances where operators want to eliminate individual channels or even the entire channel bundle they've been required to carry as a condition of gaining access to popular channels offered by a particular broadcaster, the SIPV data can help them determine the risks they're incurring in dropping channels. And it can show them which topical categories represented in a given bundle are most popular with their viewers to ensure they choose replacement channels that match those tastes.

SIPV in Real-World Deployments

The experiences of Canadian MSO Coopérative de câblodistribution de l'Arrière-Pays (CCAP) over many years of managing its bandwidth requirements through use of SIPV offers a vivid illustration of how the platform eliminates the need to spend heavily on plant upgrades or analog reclamation as market conditions require more bandwidth for DOCSIS and as TV channel counts multiply.

CCAP, which serves about 16,000 subscribers with triple-play services in the rural areas surrounding Quebec City, first implemented SIPV in 2013 in conjunction with its need to expand DOCSIS bandwidth and add more video options. Initially, with 8 U-EQAMs designated for SIPV switched video in each of thirty-six SGs, the MSO freed up sufficient bandwidth to raise its top-tier DOCSIS offering to 50 Mbps.

As a service provider enjoying over 70 percent market penetration, the motivation behind CCAP's expansion of broadband tiers and ongoing expansion of its TV channel lineup isn't so much derived from a need to gain more subscribers, but rather to prevent inroads to its base from newly aggressive competitors, including satellite providers and, especially, one of Canada's largest telcos, which has been building out fiber-to-the-home networks in the region.

Since launching SIPV CCAP has more than held its own, avoiding attrition while continuing to add TV channels to a digital lineup that now includes 392 HD and SD digital channels. Today the company is offering 142 of those channels through the SIPV platform, which means CCAP is using just 8 QAMs to deliver what normally would require 32 in traditional broadcast mode.

This is just the beginning of what CCAP has planned, as reflected in the fact that it is already planning to split several service groups and expand SIPV to an additional 16 QAMs to accommodate unlimited HD and, in the future, 4K UHD channels, while ensuring it can continue to free bandwidth for DOCSIS expansions to Gigabit (and beyond) as it moves ever more of its TV lineup into SIPV. All of this has been accomplished without investing in a forced analog reclamation.

Another benefit of SIPV is highlighted in the plans of Access Communications Co-operative Ltd., a Canadian triple-play cable operator based in Regina with over 125,000 homes passed in 230 communities throughout Saskatchewan. Access, which leverages both the ARRIS DAC (Digital Access Controller) headend platform to configure components and schedule services in house and Comcast Wholesale's HITS Quick Take for external turnkey headend support in some systems, has deployed SIPV to free up bandwidth and enable expansion of DOCSIS services quickly and cost effectively, notes Craig Van Ham, Vice President of Technology at Access.

"We currently operate both DAC and HITS-QT platforms with Motorola set-tops for video, and like most other operators, our RF spectrum is full, leaving limited room to expand our DOCSIS Internet services," Van Ham says. "We're looking forward to providing our customers with faster broadband speeds by taking advantage of the newly freed-up bandwidth."

At the same time, Van Ham notes, SIPV has made it possible for Access to accelerate headend consolidation across its territories. "Adara's SIPV solution will allow us to connect many of our current HITS- QT+ communities directly to our DAC network and video sources in Regina, enabling us to deliver more state-of-the-art video services to these towns," he says. "This will allow us to eliminate dozens of headends."

Figure 13

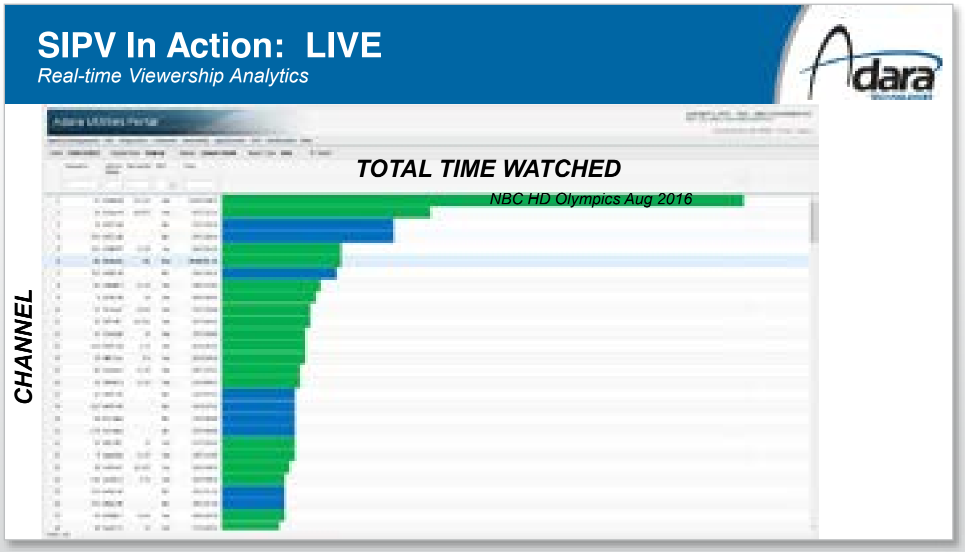

Conclusion

Cable MSOs have no margin for complacency when it comes to addressing bandwidth needs for DOCSIS expansion. Just as 100 Mbps throughput, once deemed a fantasy, has become mainstream reality, the emergence of 1 Gig services has put this new service level on track to being commonly available in the not-too-distant future. And there seems to be no prospect for letup beyond that point.

Meanwhile, TV channel expansion has gone forward at a steady space as MSOs take steps to ensure they can match or exceed offerings from satellite competitors. And it won't be long before the HD experience is replicated with the emergence of 4K UHD services from DBS and OTT outlets, which will more than double per-channel bandwidth requirements even with use of HEVC compression.

The extraordinary pace of developments calls for a better approach to bandwidth expansion than the industry has found through the costly processes entailed with plant upgrades and analog reclamation. This is why increasingly more MSOs are selecting the Adara SIPV platform as a far more cost-effective solution with an impact on bandwidth expansion that is most often superior to other approaches. In fact, by adopting SIPV as the priority strategy, they know they will be able to get far more from those other approaches if and when they become necessary.

SIPV is a game changer. It provides operators the tools they need to quickly address any immediate expansion requirements without touching outside plant or CPE. And it opens a path for ongoing expansion that will continue to serve their needs for a long time to come with little, if any, need to resort to plant upgrades. <

Footnotes

1 FCC, 2016 Measuring Broadband America Fixed Report, December 2016

2 Akamai, Fourth Quarter, 2016 State of the Internet Report, December 2016

3 Nielsen Norman Group, Nielsen's Law of Internet Bandwidth, April 1998 (updated 2016)

4 *Ibid.*

5 Deloitte Global, The Dawn of the Gigabit Internet Age, January 2016

6 Ars Technica, AT&T Names 11 New Metro Areas for Gigabit Internet, October 2016

7 Broadband & Technology News, Where in America Can You Get Gigabit Internet?, April 2016

8 Cisco Systems, Virtual Network Index, June 2016

9 Ooyala, Q2 2016 Global Video Index, August 2016

10 Yahoo Finance, Parks Associates: Adoption of Internet-Connected Entertainment Devices, September 2016

11 Broadband News, Strong Smart TV Growth in Japan, US and Europe, February 2016

12 ScreenPlays, New Initiatives Poised to Accelerate Pay TV Business Upheaval, January 2016

13 Home Media Magazine, Millennials Covet OTT Video and Pay TV, June 2016

14 Sandvine, Global Internet Phenomena Report, August 2016

15 Broadband News, Five Connected Media Devices per Home by 2019, September 2015

16 Lightreading, How Cable Plans Symmetrical Gigabit via FDX, January 2017

17 IHS Markit, TV Sets Intelligence Services,, December 2015

18 Futuresource Consulting, 4K UHD Is Key Sales Driver for Flat TV Screen Market, April 2016

19 SNL Kagan, A Clearer Picture of Growth: 2016 Global 4K UHD Industry Survey, July 2016

20 Multichannel News, Amazon Extends HDR Support to Samsung's Galaxy Note 7, August 2016

21 4K.com, Ultimate 4K Movies and UHD Content Guide, September 2016

22 Goldman Sachs, Virtual & Augmented Reality, January 2016

23 NATPE/CTA, Hollywood Sees Virtual Reality as Game Changer, January 2016

24 ScreenPlays, Getting Real about Virtual Reality, February 2016