A SOUND APPROACH TO GAINING MARKET ADVANTAGE THROUGH CPE VIRTUALIZATION

Orchestrating a New Level of Customer Experience
Requires Shrewd Use of Cloud & Premises Resources

TABLE OF CONTENTS

INTRODUCTION

Using virtualization, NSPs (network service providers) have an unprecedented opportunity to increase their market appeal across an expanding range of services and applications. This can be achieved by coordinating the expansion of functionalities between CPE (customer premises equipment), personal devices and cloud resources.

The market and strategic forces impacting the design of a new generation of STBs (set-top boxes) and Wi-Fi equipped broadband gateways and extenders require a broader perspective on vCPE (virtual CPE) development, where the goal is to implement advances in hardware processing and software design to maximum effect. Determining where functionalities should reside and creating the orchestration framework that enables resources in the cloud, CPE and personal devices to work together as a seamless whole, is mandatory for success.

By applying multi-service management capabilities to deliver quality-assured, highly secured service performance to the devices that customers connect to their premises networks, NSPs have a significant edge over a growing army of OTT (over-the-top) competitors.

Avoiding the cost-prohibitive expansion of processing power in the CPE requires judicious use of cloud software. The challenge is to define a new vCPE and premises networking architecture that can tap all available resources to maintain this advantage.

CPE virtualization is about more than saving money on equipment. vCPE solutions use cloud technology to facilitate service goals including more flexibility and faster cadence across all services, while minimizing the costs incurred with expanding CPE capabilities far beyond traditional requirements.

Whether or not virtualization technology is involved, the role of the cloud is to contribute efficiencies and functionalities that maximize operators' ability to orchestrate all elements in support of service goals. Determining where to locate processing for all the required functionalities is a matter of simple logic.

If the efficiencies of allocating hardware resources in the cloud across multiple homes or businesses can be exploited without compromising performance of a particular function, it makes sense to position the function remotely. If very low latency isn't required, the cloud becomes more attractive for that function. Otherwise, the processing belongs in the CPE.

Given the greatly expanded range of functions that must be supported to ensure NSPs maintain competitive strength and market appeal, even with the use of virtualization technology, the totality of processing tasks to be assigned to STBs, modem gateways and home networking components, will match, and may well exceed, past requirements.

Defining this virtualization-enhanced CPE and premises networking architecture starts with a clear articulation of the user experience an operator wants to deliver with each service. A fundamental goal common to all consumer and business services is to provide on-premises network connectivity over wireless and hard-wire links that meets customer expectations. Persistent throughput that is true to the subscribed broadband tier rate must be met as well as seamless continuity of experience that allows users to move from one device or one location to another without interruption. When it comes to specific services, the challenges center around the customer experience with services where the competitive threat is greatest, namely, those related to video and Internet of Things (IoT).

We begin the discussion with a look at the implications for resource allocation in NSPs' efforts to create a more robust premises networking environment where Wi-Fi connectivity has become an especially urgent area of concern. We then turn to the role expanded functionalities and their positioning locally, and in the cloud, will play in meeting NSPs' video and IoT service goals.

THE MANDATE FOR BETTER PREMISES NETWORKING PERFORMANCE

Wi-Fi has rapidly moved from being an ancillary component of premises networking to a dominant role as the number of wireless devices used to access content and applications from any point in the home or business multiplies. This trend, in tandem with demand for throughput at levels sufficient to support TV-caliber HD and even UHD video, has exposed the limits of legacy Wi-Fi solutions, setting in motion an industry-wide search for better approaches as customer dissatisfaction mounts.

In the past, Wi-Fi issues weren't regarded as an operator's responsibility. However, that attitude is rapidly evaporating, no matter what the degree of dependence on Wi-Fi as the primary networking conduit might be. Even though most NSPs will rely on wired connectivity to some degree for some

time to come, Wi-Fi has become too important to leave performance to the bring-your-own device environment.

For example, cable operators must still rely on coaxial links, typically utilizing MoCA (Multimedia over Coax Alliance) technology to transport QAM-based pay TV signals to TV screens throughout the home. As they shift to all-IP (internet protocol) video operations, they're likely to take greater advantage of the flexibility of high-performance Wi-Fi networks to transport TV signals, possibly in conjunction with hybrid platforms that tightly integrate the use of wireline and wireless connectivity in the home. Already, many telcos delivering video in IPTV mode have shifted to all-Wi-Fi networking utilizing Wi-Fi equipped STB gateways.

The Escalating Urgency for Better Performance in the Home

The urgency of the Wi-Fi issue is undeniable, as evidenced in results from a global consumer survey published by ARRIS in its 2015 Consumer Experience Index report.[i] The survey found that 63 percent of home Wi-Fi users worldwide are experiencing performance issues of one kind or another. Moreover, 72 percent of consumers said that Wi-Fi availability in every room of the home is very or vitally important, and 54 percent indicated they are not experiencing the range of coverage they need.

NSPs are experiencing a direct impact on customer care operations as a result of these issues. Some operators report that 50 percent or more of trouble calls into their call centers relate to Wi-Fi problems, even though in most instances they have nothing to do with the quality of performance on NSP-supplied gateways.

Video consumption on wireless devices has a lot to do with these issues, as does the number of devices actively used in the average household. 59 percent of consumers now watch professionally produced long-form video on mobile devices, which typically offload those streams to Wi-Fi when in the presence of Wi-Fi APs (access points).

Without NSP intervention with new approaches to Wi-Fi network management, demand for resolution of these issues is destined to intensify rapidly. The per-household wireless device count now stands at seven in North America, according to network traffic statistics compiled by Sandvine.[ii] Projections for the future vary widely. According to a study by market researcher IHS, the per-household count will climb to ten in North America by 2019.[iii] By Intel's estimation, which takes into account high expectations for IoT applications, the average will hit 50 by 2020.[iv]

As the Wi-Fi connected device count rises, the per-session bitrates tied to video usage will go up as well. Even more IP video, whether via OTT or managed IPTV connections, will be delivered by Wi-Fi to smart TVs and older TV sets connected through IP STBs as well as PCs and game consoles.

The effect will be compounded as 4K UHD enters the picture. UHD doubles the required bitrates compared to HD with use of HEVC (high efficiency video coding) or other advanced codecs. Virtual reality (VR) is also a factor. Its bitrate requirements are estimated to be in the 150 Mbps range. Already, VR has influenced market preparation for a Wi-Fi tier in the 60 GHz spectrum range with the IEEE 802.11ad standard known as WiGig.

Upstream bitrates are rising rapidly as well, with billions of high resolution camera-equipped smartphones, tablets and PCs now in use. These devices support video calling, the use of video and photos in social networking and uploads to cloud storage at transmission rates that can top 100 Mbps.

The Enterprise Perspective on the Need for Better Performance

Wi-Fi performance is a growing problem in the enterprise and institutional domains, where the same issues tied to proliferation of devices and higher bitrate consumption on those devices are in play. Video, for example, is a much bigger factor in traffic passing over business premises networks than it once was. The enterprise video impact extends to streaming or bulk file transport in marketing and training, advertising, entertainment, surveillance and video-based collaboration and conferencing.

Non-video data files, too, are bigger and more voluminous than ever. Many business and institutional sectors including health, education, finance and government are leveraging the efficiencies of big data processing and storage technology to aggregate huge volumes of data for multiple uses across multiple work stations. Growing reliance on cloud-based applications is making network connectivity for performing routine tasks the norm.

With growing dependence on Wi-Fi from APs in small offices to WLANs in larger enterprises, dissatisfaction in service performance resulting from under-performing Wi-Fi connectivity is coming back to NSPs in the commercial services arena much as it is on the residential side. A recent survey of 100 enterprise IT administrators in business and education conducted by ZK Research found that close to 60 percent of network administrators are spending more than a quarter of their time troubleshooting Wi-Fi issues, with time spent on diagnosing and resolving each issue exceeding 30 minutes.[v]

The survey found that far too often, the initial response to such problems is to blame them on the APs. While Wi-Fi malfunctions are the reported issue 64 percent of the time, issues such as device limitations, insufficient coverage or over saturation of simultaneous usage are the actual root cause in the vast majority of cases.

KEYS TO ENHANCING PERFORMANCE IN A WI-FI-CENTRIC ENVIRONMENT

There are four types of problems that typically cause a negative user experience with Wi-Fi connections in residential or business environments. These include:

1. Range limitations where the Wi-Fi signals don't reach all locations in and around the home or business.
2. Interference from other users' Wi-Fi APs operating in multi-tenant and other dense clusters, as well as from microwaves and other non-connected devices that operate in the 2.4 GHz spectrum region used by Wi-Fi.
3. The lack of roaming support in Wi-Fi technology which often results in a device "sticking" to an AP when there's another one closer at hand, causing a drain on power as the AP tries to maintain the bit flow to that device.
4. Failure of dual-band devices to take advantage of the 5 GHz channel on dual-band APs, resulting in congestion on the 2.4 GHz band.

Use of virtualization technology will play a major role in solving these issues. New functionalities supported by more intelligent Wi-Fi components on premises will also help address performance issues. To avoid customer dissatisfaction caused by sub-par Wi-Fi performance, operators could supply the necessary solutions through surcharges, inclusion of such solutions with premium tiers or as a marketing incentive across all tiers.

The good news is that NSPs can anticipate the overall costs of providing a robust Wi-Fi presence on customer premises, which need not be as high as some of the providers of proprietary solutions

vying for market acceptance suggest. Moreover, cloud-CPE orchestration creates an automated configuration environment for subscriber installation of CPE, taking the traditional tech installation costs out of the equation.

The Vital Role of 802.11ac

Before delving into how combinations of enhanced CPE intelligence and virtualization of some functions can best be applied to achieve superior Wi-Fi performance, it's important to recognize that the starting point for achieving maximum throughput on all devices is to ensure customers are equipped with APs that support the 802.11ac protocol.

802.11ac supports connectivity on the unlicensed spectrum channels authorized for Wi-Fi use at the 5 GHz tier while enabling 802.11n-compatible connectivity at the 2.4 GHz tier. While the option to use 5 GHz was incorporated into the 802.11n standard, the option has not been implemented with most 802.11n APs now in commercial use. Thus the real impact of bandwidth expansion enabled by use of 5 GHz spectrum is tied to 802.11ac.

According to ABI Research, 68 percent of all Wi-Fi devices shipped in 2015 were dual-band and 802.11ac capable, which means more dual-band than single-band devices are now in users' hands.[vi] ABI Research predicts Wi-Fi chipset shipments worldwide will total 18 billion between 2015 and the end of 2019.

802.11ac also incorporates a broad range of performance enhancements that have been incrementally implemented in three successive "waves" of chipset and OEM product releases. There is a broad disparity in performance capabilities across the three waves of 802.11ac which means NSPs, to whatever extent possible, are well advised to deploy Wave 2 or Wave 3 equipment, depending on the timing of their purchases.

Wave 1 APs began shipping in 2012. Wave 2 APs, with significant additional enhancements, began shipping in 2015 with interoperable certification by the Wi-Fi Alliance starting in June 2016. Devices and APs supporting Wave 3 are now shipping as well. However, it will be some time before the device count reaches a point where full use of the Wave 3 capabilities will be widespread.

MIMO and MU-MIMO

In all cases, 802.11ac APs have the option to employ enhancements to MIMO (multiple-input,

multiple output) technology, an optional component for both 802.11n and ac that uses multiple antenna arrays in transmitters and receivers to support more robust transmissions through spatial separation of bit segments using SDMA (space division multiple access) multiplexing.

While it hasn't been a big factor with 802.11n APs, MIMO is a mainstay in dual-band 802.11ac hardware using 20, 40 and 80 MHz channels in the 5 GHz band and 802.11n in the 2.4 GHz band. In addition, all 802.11ac devices support modulation at up to 256 QAM whereas 802.11n tops out at 64 QAM.

Wave 2 APs support MU-MIMO (multi-user MIMO) functionality. MU-MIMO uses precoding to identify multiple receiving devices, thereby enabling a kind of broadcast mode that uses SDMA to transmit content delivered over a given frequency channel to more than one device. This has the effect of doubling or even quadrupling Wi-Fi network user density by increasing the number of simultaneously supported clients.

Beamforming

Another key advance has to do with adaptive antenna technology known as beamforming, which is used to overcome the inconsistencies in quality of service caused by variations in power, types of antennas, support for IP-based streaming and other device characteristics. 802.11ac introduces a standardized approach to beamforming, eliminating the interoperability issues that plagued non-standardized implementation of beamforming on 802.11n.

In an environment where a simple variation in a device's orientation can account for up to a 5x performance degradation, beamforming provides APs with the sensitivity and functionality essential to maintaining consistent quality of service (QoS), by spatially filtering radio signals in phased arrays to create constructive interference in the direction of a specific target. This allows the transmitter to weigh phase and gain in favor of transmitter-to-receiver antenna paths that generate the highest possible signal power at the receive end, resulting in increased and more consistent overall throughput and reduced power consumption.

Performance Gradations in Succeeding Generations of 802.11ac APs

Wave 1 APs support any combination of 20, 40 and 80 MHz channels with the top level physical layer (PHY)-level throughput at 1.3 Gbps attained with five channels operating at 80 MHz. Wave 1 APs support up to three transmit and three receive antennas and are typically configured to support 2x2 configurations.

Wave 2 APs, along with supporting 4x4 configurations, use 160 MHz channels and bonding of two 80 MHz channels. With the combination of Wave 1 capabilities, MU-MIMO and 160 MHz channels, the maximum PHY rate throughput touted for Wave 2 APs is 2.34 Gbps versus 1.3 Gbps for Wave 1, which translates to maximum effective rates of about 60 percent of the PHY rates, or 800 Mbps and 1.4 Gbps, respectively.

In all 802.11ac instances, the total AP throughput includes the bitrate available over the 802.11n 2.4 GHz channels, where the maximum PHY rate is 600 Mbps. Whichever wave is implemented on the premises gateway router, 802.11ac represents a vast increase in available throughput for end users.

Sales of Wave 2-enabled mobile devices reached the hundreds of millions by late 2016, according to Strategy Analytics.[vii] The research projects that shipments of Wave 2 APs will exceed 10 percent of the market in 2017 and dominate AP sales by 2020.

Wave 3 APs will raise the maximum antenna configurations to 8x8, which with eight 160 MHz channels operating at 256 QAM raises the PHY rate to 6.933 Mbps. Realistically, given there are a limited number of 160 MHz channels available in any given region, the PHY rate and effective rate will reflect some combination of 160 MHz and narrower channels, but should at least double the effective rate compared to Wave 2 APs.

Beyond 802.11ac

Specifications for the next generation of Wi-Fi, 802.11ax, are under development with the goal of pushing the PHY rate to the 10 Gbps level. Simultaneously, the aforementioned 802.11ad WiGig protocol is being certified for introduction to the market.

The goal for WiGig is to open a path for building an ecosystem of usage in the 60 GHz tier where extraordinary levels of bandwidth are required for a given application, such as virtual reality. PHY throughput at the 60 GHz tier tops out at 4.6 Gbps.

ABI Research forecasts that by 2021 WiGig chipsets supporting tri-band usage with 802.11ac and 802.11 n, will be shipping at a rate of one billion annually with the number of tri-band smartphones in the market representing 28 percent of total smartphone penetration.[viii] 802.11ay, slated for market introduction in 2017, enhances 802.11ad by adding MU-MIMO to the GHz band, raising the PHY ceiling to 100 Gbps.

Another important protocol advance intrinsic to Wi-Fi performance enhancement is 802.11k, which is supported by growing numbers of new devices entering the market. 802.11k-equipped devices automatically switch to the strongest AP signals, obviating reliance on software in APs to switch connections, which can cause breaks in a session flow lasting as long as a few seconds. Along with automating a change in connection based on the proximity of an AP, 802.11k causes devices to connect to a more distant AP when that signal is stronger than can be obtained from an over-subscribed AP closer at hand.

However, even with implementation of 802.11ac APs and emerging enhancements introduced with 802.11k and 802.11ay, much more needs to be done to address the impediments to high performance enumerated at the outset of this section and also to enable new functions intrinsic to creating more compelling service options. The new CPE architecture must take into account the need to extend operators' management of the user experience to every device well into the future.

STEPS TO MAXIMIZING PREMISES NETWORK PERFORMANCE

Anticipating Increasing Demands on Performance

NSPs can usually achieve complete coverage of about 97 percent of subscribers' homes with no more than two 802.11ac APs, including the primary gateway router and an extender, assuming device counts are in line with the five-to-seven norm. The remaining three percent, consisting of very large homes, can nearly always be provided complete coverage with three APs.

However, device counts are likely to continue rising, especially as market adoption of IoT applications grows. Looking at IoT adoption from the device perspective, Gartner predicts the number of installed IoT-related devices in homes and businesses worldwide, not counting smartphones, tablets or PCs, will grow to 21 billion by 2020.[ix]

This represents close to a 328 percent increase over the 6.4 billion devices the firm tabulated for 2016. Gartner says that by 2020, IoT component and connectivity costs will be so low that just about everything – light fixtures, windows, doors, appliances, toys and sporting equipment – will be equipped for IoT connectivity to support control, monitoring and sensing.

The implications are far reaching. The increase in device counts raise the contention level for airtime on existing APs, leading to a need for more. More significantly, the functional requirements imposed on gateways and extenders increase as the range of applications, service features and quality assurance requirements expands.

Along with shoring up performance of Wi-Fi connectivity in this fast-changing environment, the allocation of processing resources between CPE and the cloud must also take into account the hybrid nature of the evolving premises infrastructure. While Wi-Fi is fast becoming the primary access medium on premises, there's a need to ensure wireline access and other wireless components of the infrastructure are holistically integrated into the functions performed by CPE and cloud-hosted software.

For cable operators, given their ongoing reliance on coaxial cable to support QAM-based distribution of pay TV content, there's a need to integrate the MoCA platform into the orchestrated management of connectivity. For telcos, the wireline/Wi-Fi integration mandate applies whether they're using coax, phone lines or electrical power lines to transmit their payloads.

The challenge is especially acute in the IoT domain, where the Wi-Fi-centric architecture must also accommodate operators' growing reliance on secondary and tertiary wireless connectivity over ZigBee, Z-Wave and Bluetooth links. Operators must deploy gateways, Wi-Fi extenders, STBs and even remote controls that are equipped to support short-range wireless connectivity for easy onboarding of any NSP-supplied or consumer-purchased IoT device, from simple sensors to complex appliances.

Allocating Functionalities at the Premises and in the Cloud

In this holistically managed hybrid premises networking environment, certain functions that must be supported either at the CPE level or in the cloud apply to performance across all service and device categories. Others are intrinsic to specific service and application categories.
Not only does the increase in device counts raise the contention level for airtime on existing APs, leading to a need for more APs, the functional requirements imposed on gateways and extenders increase as the range of applications, service features and quality assurance requirements expand.

Managing PHY-Level Performance

With respect to universally applicable functions, the considerations focus on ensuring consistent throughput everywhere in and around the home or business. This largely depends on controls

residing in the CPE complemented by cloud-implemented functionalities related to IP address management and Big Data analytics. Here the goal is to determine optimal placement of those controls among gateways and ancillary APs, while identifying functionalities that can be allocated to the cloud.

As for placement of intelligence and processing in the CPE, it makes sense to have most of the functionalities resident in the primary gateway. Whether used as the sole AP or in conjunction with additional extender APs, the core gateway, functioning as the RF home network controller, should be able to generate all the low-latency commands that optimize connectivity on a per-device level so that there's always just one hop between any device and an AP.

In this scenario, AP-equipped STBs are configured as ancillary devices under control of the gateway for wireless connectivity. And the primary gateway should have the intelligence to orchestrate the use of wireline and short-range wireless protocols to maximum advantage.

Going beyond the mechanisms intrinsic to 802.11ac gateways, enhanced primary gateway capabilities might include control over throughput based on bandwidth requirements of specific devices for specific services. For example, when premium video is delivered to managed clients, STBs might be accorded a higher share of available bandwidth than video delivered to unmanaged clients like smartphones and tablets, while OTT video, best-effort data, smart home apps and other categories might be assigned different shares of capacity.

In the future, it's likely the RF network controller functionality will have to include a means by which spectrum band utilization is assigned on a service-specific basis, bringing bands beyond the two utilized with 802.11ac. For example, 802.11ah operating in the 900 MHz band provides an ideal solution for delivering signals related to low bitrate, low-power IoT applications at longer distances, even as 802.15.4 (ZigBee) and Bluetooth Low Energy (BLE) are used for short-range connectivity.

The ability to assign 60 GHz 802.11ay (WiGig) connectivity to specific applications for VR is likely to occur. There may be other uses for 802.11ay that the gateway RF controller will need to assign, such as leveraging the protocol's high throughput to provide backbone connectivity for the growing number of APs.

Whatever spectrum tiers are in use, the gateway can leverage a hub-spoke configuration with the addition of extender APs to perform the steering functions that connect client devices to the best signal source and ensure that extenders deliver the throughput optimized to device and service requirements. While extenders must have enough processing power to act on these commands,

this configuration ensures the costs of adding extenders are minimized without sacrificing overall network performance.

This approach to premises network orchestration positions the gateway as the data collection point that can provide visibility into all network components in support of cloud-hosted diagnostics functions. Client steering and other processes depend on real-time visibility into what's happening with each connection through data exchanges between the gateway and extenders.

Through automated collection, aggregation and analysis of this data in the cloud, operators can track what's happening over time in every household and business. Via automated feedback from such analysis, the RF control mechanisms in the CPE can be reset when the performance is deemed to be subpar.

The Need to Achieve Direct Control over Device QoS

The ability to perform per-device diagnostics to assess the QoS on connectivity to any CPE element is essential to both proactive customer care, and responsive care provided by CSRs or via self-help portals. The data processed by an analytics engine in the cloud be used to help subscribers mitigate issues. For example, this data could guide them in placement of extenders. Or by identifying device issues, such as when a dual-mode device capable of using both the 2.4 and 5 GHz spectrum tiers fails to jump to 5 GHz when the traffic over the 2.4 GHz channel is slowing things down.

There are multiple other uses of cloud-based analytics tied to data collection from premises gateways. A back-office system using this data can provide technicians, engineers, operations managers and senior management with performance data formulated for their specific needs from each premises footprint at whatever levels of aggregation are useful to executing their responsibilities.

Currently, cable operators who rely exclusively on DOCSIS OSS are at a disadvantage when it comes to enabling the range of customer care and administrative support functions discussed here. While DOCSIS OSS has made SNMP (Simple Network Management Protocol) the cornerstone for supporting management of fault, configuration, accounting, performance and security applications that are essential to running a high-speed data service, SNMP was not designed to meet the needs of a marketplace where subscribers expect to be able to access services via any network-enabled device they happen to buy.

With reliance on the gateway as the point of OSS interface with devices privately addressed via the NAT (Network Address Translation) process enabled by DHCP (Dynamic Host Configuration Protocol) there's no way for the SNMP system servers to directly initiate connections and management functions on those devices. Firewalls, too, can present a barrier to getting SNMP into the local area networking environment.

Over time, the intermediary role of the gateway for OSS management will be eliminated with the prevalent use of IPv6. With its use, operators will be able to directly manage all devices. This will bring greater utilization of the cloud, from the provisioning of addresses to orchestration of all resources in support of service goals.

As long as DHCP remains essential to IP address management, there are new hardware processing requirements that need to be implemented at both the premises gateway and remotely, to give operators greater diagnostic visibility into the entire premises device ecosystem. Fortunately, cloud technology makes it possible to offload DHCP as a gateway processing requirement, which will partially compensate for the additional processing required for the advanced CPE management architecture discussed here.

Leveraging TR-069 and DOCSIS OSS

There are proprietary alternatives to SNMP-based device management in the market. Yet, the most practical option for NSPs is to adopt the standardized, market-proven approach embodied in the set of protocols developed through the Broadband Forum's Technical Report-069 initiative, more formally known as the CPE wide area network management protocol.

First published in 2004, TR-069 until recently was primarily used by telephone companies to manage STBs, gateways and other devices on consumer premises. Over the past few years, the protocol has rapidly moved into the cable broadband market with implementation on gateways deployed by many leading MSOs. Notably, TR-069 has been designated as an option through APIs (application program interfaces) incorporated into the RDK (reference design kit) protocol stack spearheaded by Comcast, Time Warner Cable and Liberty Global as a template for next-generation CPE.

Now in use on an estimated 356 million devices worldwide,[x] TR-069 is a bidirectional Web-services protocol using XML (extensible markup language) for its message format that defines an application layer for remote management of end-user devices employing tools based on SOAP (simple object access protocol) and IP connectivity based on HTTP (hypertext transfer protocol). Device

management in this client/server architecture is under the control of an ACS (auto configuration server) located in the service provider cloud. The ACS can manage devices that use NAT addresses and that sit behind firewalls.

In another critical departure from how SNMP works, TR-069 allows connections and interactions to be established by either the device or the ACS, without requiring persistent connectivity. Once connection is established, the ACS discovers capabilities of individual devices and supports vendor-specific parameters to enable a wide range of applications and service categories. Functionalities include auto-configuration and dynamic service provisioning, software/firmware image management, software module management and data collection from each device for use in diagnostics and status, and performance monitoring.

Bandwidth management and QoS support are vital components as well. The protocol and its extensions provide the means for classifying, policing, shaping, tagging, queuing and scheduling traffic whether connectivity is direct from the network to a device or is implemented on devices managed by a gateway. This includes the ability to extend TR-069 capabilities to devices not equipped to serve as TR-069 end points through proxy management of other protocols that might be native to those devices, such as UPnP (universal plug & play), DLNA (Digital Living Network Alliance) and OMA-DM (Open Mobile Alliance Device Management).

Through its multiple extensions, TR-069 capabilities have been expanded to facilitate the operational environments of specific device types including versioning models, use profiles and data models for broadband gateways, VoIP CPE, femto access points, PON devices, storage service devices and Wi-Fi APs. The Wi-Fi gateway model supported by the extension known as TR-181 has been made an integral part of CableLabs' Wireless Specifications suite for managing the Wi-Fi air interface in residential, enterprise and public deployments.

Importantly for cable operators, it's possible to coordinate the OSS capabilities of DOCSIS and TR-069. DOCSIS-enabled gateways communicating with NAT-addressed devices can make it possible for any DOCSIS-compliant device that also is equipped with TR-069 functionality to generate non-DOCSIS data flows by pinging the network for an ACS connection.

OPERATIONAL ENHANCEMENTS TIED TO SPECIFIC SERVICE APPLICATIONS

Video and the Set-Top Box

Nowhere is cloud technology more important to service flexibility, richer user experience and speed to market than in the pay TV realm, where virtualization of the STB has been gaining steam rapidly in NSP initiatives worldwide. NSPs understand this is necessary in order to seize the competitive advantage that lies with offering subscribers greater packaging flexibility, access to third-party OTT video options, personalization of the user experience with consistency across all devices and new levels of monetization through targeted advertising.

The opportunity to elevate pay TV to must-have status for today's "cord never" generation and an increasingly restless population of existing subscribers through development of a personalized, feature-rich experience that is far superior to the legacy pay TV experience is underscored by a recent report from 451 Research. In a global survey of 2,000 consumers, 78 percent of respondents (86 percent of Millennials) said they would be more willing to maintain their current pay TV subscription if it offered a single source to easily search, discover and watch content, including the ability to watch online video and access content from OTT services.[xi]

The survey also revealed how important personalizing such experiences can be. Nearly half of the respondents across all age brackets and 54 percent of Millennials indicated they expect their pay TV providers to use the data they collect to provide custom-tailored service offerings and customer support.

Leveraging the Cloud to Deliver a New Customer Experience

Such sentiments suggest NSPs who can enable holistic navigation on a UI (user interface) that personalizes pay TV and OTT options have a great advantage over OTT players who are incapable of providing a universal navigation experience. By pulling user data to the cloud and applying advanced analytics to assess individual preferences, operators can deliver recommendation elements and targeted advertising and in-house upsell offers suited to individual tastes over each user's UI.

Such measures build on an already significant shift of traditional set-top functions to the cloud as reflected in widespread implementation of cloud DVR (digital video recorder), which Parks Associates projects will be available to 24 million pay TV subscribers by 2018.[xii] Recent research from U.K.-based Technavio projects the cloud DVR market will grow at a CAGR of more than 30 percent through 2020, representing a growth in network and satellite operator spending from $1.48 billion globally in 2015 to $5.51 billion in 2020.[xiii]

At the same time, growing numbers of NSPs are taking advantage of the cloud to support other advanced features tied to blending of legacy and OTT services, personalized navigation and advertising. This is evidenced in widespread adoption of the CloudTV platform developed by ActiveVideo, a pioneer in pay TV cloud technology jointly owned by ARRIS and Charter Communications.

With the ActiveVideo platform, NSPs, including some of the largest cable MSOs in the world, are able to deliver UIs and applications developed in HTML5 along with user-selected viewing options over either traditional pay TV channels or broadband IP, depending on the type of device used. A thin software client in the device is used to send key clicks from standard remote controls to the cloud system, which determines which actions are required and controls all systems that are involved in executing those actions.

The cloud technology makes it possible for NSPs to deliver a next-generation user experience to legacy set-tops over MPEG QAM channels as well as to any IP-connected device in formats suited to each viewing screen at speeds comparable to set-top-based navigation systems. UIs are composed out of picture elements, that have been individually encoded to ensure bandwidth is allocated in proportion to what is required by each element, at optimal window size and frame rates for each element of the screen—and taking into account whether the session is to be distributed via the VBR (variable bitrate) mode of multiplexing used with MPEG-2 channel streams or the IP fragmentation process used with ABR (adaptive bitrate) streams.

In instances where the pay TV service is delivered in IP mode, cloud technology allows operators to shift the STB tuning function to the cloud in edge locations where latency does not become an issue as the all-IP user base scales. The process in the IP domain is not actual channel tuning, instead it's service selection whereby the cloud software formats and streams the chosen programming to the user's device over whatever type of connection is in use, be it HFC, fiber, Wi-Fi or cellular.

The cloud also has much to offer when it comes to advertising, again as evidenced by how some operators are using the ActiveVideo technology. The firm's CloudTV AdCast system supports dynamic ad insertion in the video streams or placement of banners or other types of promotions in

the UI, enabling device-independent delivery of targeted local and national TV advertising at scale. While relying on HTML5 at the cloud-level of ad selection, AdCast works whether or not the end device is browser enabled, so the addressable ads reach users on traditional set-tops, commercial IP set-tops and connected TVs that may be using specialized under-powered browsers.

Taking advantage of cloud technology also streamlines service set up. Rather than sending technicians, operators can send subscribers packaged software that will allow terminals to configure themselves on the network and to extend configuration to all other devices in accord with the policies appropriate to each one.

New CPE Processing Requirements

The use of virtualization technology seems essential to enabling a next-generation multiscreen TV service. While there are significant savings to be realized in operations and customer care costs, virtualization of set-top functions should not be viewed as a way to lower the capital costs of delivering pay TV.

Not only does the cloud processing discussed in this paper require new spending, but the tasks that must remain the responsibility of CPE bring new processing costs into the equation even as old ones are mitigated to some degree by the cloud. It's a tradeoff that leaves the CPE memory and MIPS (million instruction per second) quotients unchanged. Cases in point include tasks related to decoding, rendering, securing and caching content.

For example, the emergence of UHD 4K mandates that new set-tops be equipped to perform decoding on HEVC compressed content, representing a significant increase in processing requirements over MPEG-2 and MPEG-4 AVC (advanced video coding). STBs will also need to be able to support refresh downloads in the event of decoder failures.

While cloud-based, personalized UI construction and delivery to each user over IP or MPEG transport feeds can eliminate the need for a graphic processer unit, with the graphics-rendering power in the STB for HD and SD UIs, operators are likely to determine that the much higher throughput required for transmitting a fully rendered UHD-caliber UI from the cloud incurs too high a cost in bandwidth. In that case, the primary media gateway in the home would have to be able to handle the rich graphics rendering required with such UIs. Spending for such capabilities on ancillary STBs can be avoided, assuming the media gateway is able to transmit the fully rendered UI over the home network.

While the removal of the DVR from the STB eliminates the traditional personal recording processing and storage requirements, storage processing supporting IP video caching akin to CDN (content delivery network) functionality is destined to become an important feature of the primary STB. Such capabilities will be required for short-term storage tasks suited to enabling trick-play applications and catch-up viewing on live as well as VOD (video on demand) content.

More broadly, caching at the premises level helps operators minimize the external network bandwidth consumption that comes with delivering each piece of content over unicast streams to anyone viewing that content in a given household. Along with the caching functionalities, this requires that the media gateway be equipped to support ABR (average bit rate) streaming to every device in the home.

As STBs are relieved of the cable card burden with implementation of DCAS (downloadable conditional access system), new requirements tied to securing high-value content will impose new functionalities on STBs. The need for new mechanisms is reflected in the ECP (enhanced content Protection) specifications issued by the motion picture studios' tech consortium MovieLabs, which are widely viewed as a template for all types of premium content as producers take action against mounting levels of piracy worldwide.

Vendor implementations of server-based injections of forensic watermarks on a per-session basis have removed the watermarking process required by the ECP specs from the STB, and many other aspects of the MovieLabs requirements are already common practice in the pay TV business. However, one of the ECP requirements associated with what is called "Secure Media Pipeline" (SMP) requires a different CPU (central processing unit) architecture in the STB.

While NSPs leveraging state-of-the-art content protection technology are well in line with most of SMP requirements for "end-to-end protection that encompasses, as a minimum, decryption through to protected output," the specs add a new element by stipulating that this protection must include support for "a secure processing environment isolated by hardware mechanisms running only authenticated code for performing critical operations." In other words, at no point in the delivery chain can the CPU used to support other processes be involved in supporting access to decrypted content.

In this new pipeline environment, the series of scrambling and descrambling steps from the point of reception to screen display are serially coupled to the security system and implemented on hardware separate from the CPUs. This is a stringent requirement that disqualifies use of older STBs where a single CPU does everything.

Moreover, STBs supporting delivery of both legacy pay TV channels and IP streams to TV sets will have to be equipped to handle DRM (digital rights management) processing. ECP specs significantly raise the bar on DRM-related protection mechanisms by requiring that all IP devices qualified to receive premium content be equipped with "factory burned" hardware roots of trust.

RoT (roots of trust) have long been fundamental to STB support for CA (conditional access) systems. Today, hardware RoT will have to be implemented on STB chips at the factory to "provide a secure mechanism for DRM systems to store secrets in local, persistent storage in a form encrypted uniquely for the device," as described by the ECP specs. This stipulation is already impacting OEM suppliers of personal devices, frequently resulting in use of chipsets with embedded RoT based on the TrustZone architecture developed by mobile chipmaker ARM.

It's also important to note that there are other emerging security requirements that could impact STB designs. Notably, DECE (Digital Entertainment Content Ecosystem), the studio-backed consortium spearheading the UltraViolet digital locker initiative, and SCSA (Secure Content Storage Association), the download-to-personal-storage service developed by a coalition of studios and electronics manufacturers, have developed security measures specific to their business models that will have to be followed by NSPs who provide support for either or both services to their subscribers.

Internet of Things Services

Researchers continue to project massive growth for IoT applications with some reports pointing to recent surges in penetration as justification for the exuberance. For example, IHS reports that 15 percent of North American households had at least one smart home device, excluding devices associated with professional security monitoring, with a total of 52 million devices installed as of year-end 2015. This increased from 10 percent the year before.

IHS anticipates the penetration will hit 34 percent with 900 million installed devices by 2025.When combined with professionally monitored smart homes, including security services provided by NSPs, the total smart home penetration will reach 53 percent by 2025.

The Impact of Unresolved Issues

Uncertainty has entered the market with some researchers offering more pessimistic projections based on current readings of market penetration. These projections are well below that of IHS and

even farther off the latest figures from Parks Associates, which puts smart home device penetration at 20 percent of U.S. households. For example, Forrester says U.S. household penetration stood at 6 percent as 2016 came to a close and predicted it would only reach 15 percent by 2020.[xiv] Argus Insights reported in mid-2015 that its research showed consumer demand for IoT devices, based on expressions of intent to purchase, actually dropped by 15 percent from a year earlier.[xv]

Things look more solid on the business side. Global enterprise spending on IoT products and services was pegged at $120 billion in 2016 by IoT consultancy Zinnov, which projects the total will grow to $253 billion by 2021.[vi]

The biggest segment in the enterprise space, industrial IoT, once known as machine-to-machine involving the integration of machinery with industrial networks, has been growing rapidly for years, with the global market now valued at close to $100 billion on its way to $151 billion in 2020, according to researchers MarketsandMarkets.[xvii] However, IoT growth in the SMB market has been much more limited with 14 percent of businesses in this category now estimated to be using such products and services, according to Forrester.[xviii]

As for the non-industrial consumer and business IoT services arena there needs to be a significant course adjustment if the potential is to be realized. On the NSP side, there are rising concerns among operators that the traditional turnkey licensed-software approach falls short from a cost/benefit perspective, restricting their opportunities as consumers respond to a flood of devices and systems on offer from retail and OTT outlets: supporting controls over lighting, window shades, garage doors, appliances, home media components, sprinkler systems, smoke detectors and much more.

Things aren't going all that well in the OTT IoT domain where big gambles like Google's $3.2 billion buyout of connected thermostat maker Nest and other acquisitions such as home security camera maker Dropcam and home hub maker Resolv, since shuttered, have not produced anticipated results. Nest sold just 1.3 million smart thermostats in 2015 compared to 1.2 million the year before, according to Strategy Analytics.[xix]

Many issues stand in the way of greater success in the residential market, starting with the fact that buying application-specific products with separate management systems becomes less appealing as the number of devices and applications multiply. The more complicated any product or service turns out to be, the greater the damper on accumulating more.

Reliability and lack of tech support are big issues. In a 2014 study, Parks Associates reported that 25 percent of U.S. broadband households with smart home devices experience problems on a monthly

basis. The study found that 57 percent of consumers with an IP security camera and 58 percent using smart electronic door locks or thermostats would welcome tech support services for those solutions.

As hacking and the prevalence of malware become ever greater concerns in the public consciousness, security is becoming another problem for consumers in the IoT space. And for good reason. In one of many recent reports attesting to high levels of vulnerability, HP's Fortify on Demand unit says tests of ten leading retail home security systems found all of them to be woefully lacking in basic protections against hackers.[xxi]

This study followed another more general HP report that looked at 10 devices across multiple IoT product types. In that research the firm found an average of 20 vulnerabilities per system, spanning TVs, thermostats, home automation hubs, alarm systems, etc. The researchers saw credentials being sent over clear text, network ports listening with root shells without a password, and private data leakage as well as vulnerabilities common to Web and mobile usage in general.

The Improved Prospects for Device and Platform Interoperability

These problems have inspired new initiatives among major players both at the individual vendor level and with broader consensus on standards. Where standards are concerned, a major advance toward enabling interoperability across the IoT ecosystem was accomplished in October 2016 with the merger of the OCF (Open Connectivity Foundation) and the AllSeen Alliance, backers, respectively, of the IoTivity and AllJoyn open source projects.

Reflecting the scope of the industry's commitment to eliminating the market confusion stemming from the clutter of proprietary solutions, the expanded OCF board of directors includes executives from a wide array of leading companies including AB Electrolux, Arçelik A.S., ARRIS, CableLabs, Canon, Inc., Cisco Systems, Inc., GE Digital, Haier, Intel, LG Electronics, Inc., Microsoft, Qualcomm, Samsung Electronics and Technicolor SA.

By combining the best of IoTivity and AllJoyn, OCF offers a rich feature set for both onboarding and security coupled with a set of supported cloud protocols that range across residential, smart city and M2M industrial applications. As a result, the specifications generated by OCF will ensure that billions of connected devices from phones and computers to gateways, hubs and sensors will be able to communicate with one another regardless of manufacturer, operating system, chipset or physical transport.

The NSP Advantage Amid a New Wave of OTT Initiatives

Meanwhile, leading players in the OTT IoT space are forging ahead with new initiatives aimed at drawing greater consumer engagement. The reorganization of Google under the Alphabet corporate umbrella, with Nest as a separate operating entity, has brought a shift in strategy where the emphasis is now on what Nest calls the "Works with Nest Platform," an interoperable ecosystem open to bringing in devices and service support from partners such as ADT, Lutron, Philips, Control4, Vivint Home, Comcast, AT&T Digital Life and Amazon's Echo.

Amazon's Echo strategy is another case in point where the platform, in this case a smart hub with advanced voice recognition capabilities, has been opened to outside developers. Following suit, Google recently introduced the Google Home hub and Apple enhanced its HomeKit software suite to enable use of the iPad as a similar type of hub.

These types of steps have raised the stakes for NSPs who are faced with the decision to move ahead with a new approach to IoT services that expands consumers' options while unburdening them of the costs and hassles that come with setting up multiple, incompatible solutions throughout the home. Leveraging the interoperability promised by OCF, operators are much better positioned than OTT suppliers to bring IoT applications into the mainstream by virtue of the higher levels of convenience and performance they can deliver in a multi-app managed service environment.

NSPs have end-to-end management control over the networks that deliver these services. Additionally, they have many other assets to leverage, including: technicians in the field, direct customer relationships, the ability to amass and put to good use immense amounts of data, and a vendor supply chain offering a new generation of solutions for building holistic, multi-application services. With a horizontally extensible operations template for enabling execution across myriad verticals, operators can lower their risks in what has been a low-ARPU market segment that depends on high volume to drive ROI.

By implementing highly scalable solutions meeting current requirements with a good path forward to more applications and features, NSPs will find their advantage is strengthened as new applications emerge to capture consumer interest. While operators can't support everything consumers might buy from local stores, they can identify certain brands of devices and applications that customers can plug into their smart home service without going through the installation and application launch hassles associated with each individual device.

Comcast's Xfinity Home exemplifies this approach, where an open source-based platform delivers integrated home security, automation and energy management solutions that allow users to monitor and

control smart home functions from one app, touchscreen or web portal. Through its Works With Xfinity Home program, Comcast has enabled subscribers to readily identify and select devices offered at retail that will work on the platform, such as the Nest Learning Thermostat, Chamberlain MyQ garage controller, Lutron Caséta wireless controller and dimmer and August Smart Lock.

The Cloud Role in NSP IoT Strategies

Obviously, the cloud has had a huge role in IoT development with support for device on-boarding, application provisioning and management, device and user authentication, analytics and more. By deploying a versatile cloud-based software control plane that's designed to interface with existing middleware, back-office components and new service support modules, operators can bring a virtually unlimited range of options into the user experience.

Utilizing a wide range of open APIs, NSPs can orchestrate the management of those options, usage policies, and how those options are exposed on a personalized basis across the whole community of users. With advanced analytics tools, operators can support multiple dynamic business models across complex IoT value chains associated with a wide variety of services and devices. This includes providing frameworks for accelerating growth of the operator's connected device ecosystem by empowering partners to design, launch and manage their own connected devices on the network.

Moreover, operators can formulate their management platforms to ensure IoT messaging and controls are available on smartphones, tablets, PCs, STBs and smart TVs to fully integrate IoT applications into their full service offerings. Again, the goal is to have in place a service platform that can accommodate virtually any application with data-gathering, personalization and ease-of-use capabilities that far outperform what consumers can obtain going to OTT or traditional retail sources for piecemeal solutions.

The IoT-Optimized CPE Ecosystem

With all the functionality the cloud can deliver, there needs to be a CPE ecosystem equipped with processing power to support applications and orchestrate connectivity across all wireless and wireline links. Separate IoT hubs or controllers embedded in the primary gateway must be able to manage personal app devices, cameras and sensors, and all the data flowing from those devices in accord with the state of each user. Moreover, operators will need to deploy gateways, Wi-Fi extenders and STBs that are equipped to support short-range wireless connectivity protocols like IEEE 802.15.4, used with ZigBee connections, and BLE.

There are less obvious requirements that must be accommodated by CPE-based intelligence. For example, with mounting volumes of data flowing back and forth from the cloud to orchestrate provisioning, personalization, feature upgrades and performance data collection, operators will likely want to use compression technology to save bandwidth, which means the transmitting IoT hubs must be able to support encoding and decoding of the data.

Another area of functionality that will require additional CPU power in core CPE involves support for securing data flows. For example, in the case of protecting against hacker intrusion into surveillance, healthcare or other systems packed with sensitive information, the hub may be tasked with matching all outflowing video and other data to the originating devices to determine whether the transmission carries security-sensitive information. If it does, the hub would then need to read the IP packet headers to ensure the information is directed to a legitimate destination and to block it if it isn't. In the IoT domain, as in other areas where CPE virtualization is present, the memory and CPU requirements will be substantial. A case in point is the gateway Comcast uses with the Xfinity Home service, where many of the hardware-based functionalities discussed here have been implemented.

Due to the progress toward standardization represented by the expanded OCF initiative, such gateways can now leverage the APIs developed through IoTivity or AllJoyn to support the appropriate type of connectivity for any given device and application, including ZigBee, BLE, and Google's IPv6-based Thread as well as Wi-Fi spectrum tiers. With such components in place, NSPs can provide users a broad choice of devices which they can be assured will be able to connect through a single point of communication to the cloud.

CONCLUSION

The onset of CPE virtualization represents an entirely new approach to developing and supporting a next generation of NSP video and IoT services. With the flexibility to address customer demand in ways that deliver the services and applications people want on virtually any device they choose, operators have an opportunity to create customer experiences with those services and apps that can't be duplicated by OTT competitors.

Defining the vCPE and supporting cloud architecture that achieves these ends requires a holistic view of all the hardware and software resources that will ensure that functionalities are positioned for optimum results. The cloud can be used to support any processes where the scalability of reach

and development flexibility can be exploited without any compromise in performance. vCPE, with some traditional as well as new functions offloaded to the cloud, can be equipped to perform a multitude of low-latency and other functions new and old that are essential to realizing the new service goals.

The ability to orchestrate cloud and CPE operations as a seamless whole enables use of analytics and other back-office resources to support personalization essential to compelling navigation and discovery, feature-rich applications and new approaches to monetization – from versatile packaging to addressable advertising. Equally important, orchestration of cloud and CPE functionalities makes it possible to support an increasingly wireless-centric premises networking environment that delivers the quality of performance essential to meeting customer expectations on all devices across every square foot of the home or office, including exterior as well as interior spaces.

RESOURCES

v. ARRIS, Consumer Entertainment Index, July 2015
vi. Sandvine, Global Internet Phenomena Report, August 2016
vii. Broadband News, Five Connected Media Devices per Home by 2019, September 2015
viii. The Economic Times, Households Have 10 Devices Now, Will Rise to 50 by 2020, August 2016
ix. ZK Research, IT Administrators Struggle with WiFi Troubleshooting, October 2016
x. ABI Research, 5 GHz Wi-Fi Is Now Mainstream, April 2015
xi. Strategy Analytics, 802.11ac Wave 2 with MU-MIMO: The Next Mainstream Wi-Fi Standard, November 2015
xii. ABI Research, Wi-Fi Quarterly Report, November 2016
xiii. Gartner, Gartner Says 6.4 Billion Connected Things Will Be in Use in 2016, November 2015
xiv. Broadband Forum, Ovum Survey Results, June 2015
xv. 451 Research Group, What Do Pay-TV Customers Really Want? May 2016
xvi. Parks Associates, press release, March 2015
xvii. Technavio, Global Cloud DVR Market 2016-2020, June 2016
xviii. The Economist, Connected Homes Will Take Longer to Materialise than Expected, June 2016
xix. Argus Insights, Smart Home Report, July 2015
xx. Zinnov, Internet of Things Technology Services, August 2016
xxi. Markets and Markets, Industrial IoT Market by Technology, January 2016
xxii. Forrester, The Internet of Things Heat Map, January 2016
xxiii. The Economist, *ibid*
xxiv. Parks Associates, Evolution of Smart Home & the Internet of Things, August 2014
xxv. HPE, IoT Is the Frankenbeast of Internet Security, February 2015