ARRIS

THE IMPLICATIONS OF CPE VIRTUALIZATION

Putting cloud technology to work to support next-generation subscriber services

TABLE OF CONTENTS

PUTTING THE CLOUD TO WORK

The onset of customer premises equipment (CPE) virtualization represents an entirely new approach to developing and supporting a new generation of managed network video and Internet of Things (IoT) services.

Service Providers will now have the flexibility to address customer demand in ways that deliver the services and applications people want on virtually any device they choose. They'll also have an opportunity to create customer experiences that can't be duplicated by over-the-top (OTT) competitors. For all this to happen, a new way of thinking about where functionalities should reside is required. The result of this mindset shift has led to the creation of orchestration frameworks that enable resources in the cloud, CPE and personal devices to work together as a seamless whole.

Seen in this light, CPE virtualization is not about saving money on equipment; it's about putting cloud technology to work where it facilitates achievement of service goals. These goals include more flexibility and faster cadence across all services, while limiting costs incurred with the expansion of CPE capabilities far beyond traditional requirements.

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.

Adhering to this principle results in allocation of many tasks to the cloud where minimal latency and other requirements of functions best left to CPE are not an issue. Service Providers will need to greatly expand their range of functions to ensure they maintain competitive strength and market appeal. Given that need, even with the use of virtualization technology, such as virtualized CPE (vCPE), the totality of processing tasks that will have to be assigned to set-top boxes (STBs), Wi-Fi gateways and other CPE will certainly match, and may well exceed, past requirements.

STEPS TO MAXIMIZING ON-PREMISES WI-FI PERFORMANCE

Wi-Fi has rapidly moved from an ancillary component of premises networking to a dominant one 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.

Operators have gained some relief from these performance pressures with widespread installation of 802.11ac Wi-Fi gateways. And, over time, they will be able to begin using gateways that support more advanced capabilities.

For example, 802.11ax will push the physical interface (PHY) rate to the 10 Gbps level. Closer at hand, 802.11ad (WiGig) will open a path for building an ecosystem of usage in the 60 GHz tier where PHY throughput over spectrum open to Wi-Fi will max out at 4.6 Gbps.

However, even with these and other advances taking shape on the horizon, there's much more that will need to be done to both address impediments to high performance and to enable new functions intrinsic to creating more compelling service options. The new CPE architecture must take into account the need to extend Service Providers' management of the user experience to every device well into the future.

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. In other words, while Wi-Fi is fast becoming the primary access medium on premises, there's a need to ensure MoCA® technology and other types of wireline access, as well as non-Wi-Fi wireless components, are holistically integrated into the functions performed by CPE and cloud-hosted software.

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

Managing PHY-level performance

Not only does the increase in device counts raise the contention level for airtime on existing access points (APs), leading to a need for more APs, but the functional requirements imposed on gateways and extenders increase as the range of applications, service features and quality assurance requirements expand as well.

Consistent throughput everywhere in and around the home or business depends on controls residing in the CPE. The CPE is complemented by cloud-implemented functionalities related to Internet Protocol (IP) address management and big data analytics. The challenge is to determine the 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 reside in the primary gateway. Whether used as the sole AP or in conjunction with additional extender APs, the core gateway, functioning as the radio frequency (RF) home network controller, should be able to generate all the low-latency commands that optimize connectivity on a per-device level. That way 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 when it comes to wireless connectivity. And the primary gateway should have the intelligence to orchestrate the use of wireline and short-range wireless protocols to maximum advantage as well.

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, premium video delivered to STBs might be accorded a higher share of available bandwidth than video delivered to unmanaged clients like smartphones and tablets. At the same time, OTT video, best-effort data, smart home applications and other categories might still be assigned different shares of capacity.

Over the long run 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 into play bands beyond the two utilized with 802.11ac, such as WiGig and 802.11ah. The latter, 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.

Whatever spectrum tiers are in use, the gateway can leverage a hub-spoke configuration with the addition of extender APs. The configuration allows steering functions that connect client devices to the best signal source and ensure that extenders deliver the throughput optimized for device and service requirements. While extenders must have enough processing power to act on these commands, this configuration ensures the cost of adding extenders is minimized without sacrificing overall network performance.

The need to achieve direct control over device QoS

The hub-spoke approach to premises network orchestration positions the gateway as the data collection point, providing 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, Service Providers can track what's happening over time in every household and business. Via automated feedback from the cloud, the RF control mechanisms in the CPE can be reset when the performance is deemed to be subpar.

The ability to perform per-device diagnostics to assess the Quality of Service (QoS) of connectivity to any CPE element is essential to both proactive customer care where problems are fixed before they generate complaints, and responsive care, as provided by customer service representatives or via self-help portals.

These functionalities require more granular per-device visibility and control than can be accomplished through sole reliance on Simple Network Management Protocol (SNMP)-based DOCSIS OSS.

DOCSIS OSS has made SNMP the cornerstone for supporting management of fault, configuration, accounting, performance and security applications that are essential to running a high-speed data service. However, SNMP was not designed to meet the needs of a marketplace where subscribers expect to be able to access services via the network-enabled device of their choice.

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

Consequently, there are new hardware processing requirements that need to be implemented at both the premises gateway and remotely to give Service Providers greater control over, and diagnostic visibility into, the entire premises device ecosystem. This requires use of the remote per-device management capabilities of the Broadband Forum's Technical Report-069 protocol and its extensions.

This need has been recognized with the inclusion of TR-069 in ARRIS and other vendor gateways and as an option in the Reference Design Kit (RDK) protocol stack spearheaded by leading pay TV providers. More than 40 million devices running RDK have been deployed by Service Providers worldwide. TR-069[1] is a bidirectional web services protocol using Extensible Markup Language (XML) for its message format that defines an application layer for remote management of end user devices employing tools based on Simple Object Access Protocol (SOAP) and IP connectivity based on Hypertext Transfer Protocol (HTTP).

Device management in this client/server architecture is under the control of an auto configuration server (ACS) located in the Service Provider cloud. The ACS can manage devices that use NAT addresses and that sit behind firewalls. DOCSIS-enabled gateways communicating with NAT-addressed devices can make it possible for any DOCSIS-compliant device that is also equipped with TR-069 functionality, to generate non-DOCSIS data flows by pinging the network for an ACS connection.

TR-069 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. The protocol and its extensions also 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 a device managed by a gateway.

Moreover, TR-069 extensions have been expanded to facilitate the operational environments of specific device types, including versioning models, use profiles and data models for broad-band gateways, VoIP CPE, 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. These specifications allow for the management of Wi-Fi air interface in residential, enterprise and public deployments.[2]

OPERATIONAL ENHANCEMENTS IMPACTING VIDEO AND THE STB

Virtualization is also playing a major role in the evolution of the STB. By pulling STB-generated data to the cloud, Service Providers can apply advanced analytics to assess individual customer preferences. With these detailed insights, Service Providers can pursue the competitive advantage that comes with offering subscribers greater packaging flexibility, access to third-party OTT video options and personalization of the user experience with consistency across all devices. This advanced type of STB can also facilitate monetization through targeted advertising.

Proven benefits of cloud-augmented STB capabilities

Cloud-augmented STB capabilities may be deployed via the CloudTV platform developed by ActiveVideo, a pioneer in pay TV cloud technology jointly owned by ARRIS and Charter Communications. Service Providers utilizing the ActiveVideo platform, including some of the largest cable MSOs in the world, are able to deliver user interfaces (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. [3]

Utilizing a thin software client in the STB that sends remote control commands to the cloud system, Service Providers are able to deliver a next-generation user experience to legacy set-tops over MPEG QAM channels. The same can also be done for any IP-connected device in formats suited to each viewing screen at speeds comparable to STB-based navigation systems. UIs are composed of picture elements that have been individually encoded to ensure bandwidth is allocated in proportion to support the optimal window size and frame rates for each element of the screen. This takes into account whether the session is to be distributed via the variable bitrate (VBR) mode of multiplexing used with MPEG-2 or MPEG-4 AVC channel streams, or the IP fragmentation process used with adaptive bitrate (ABR) streams.

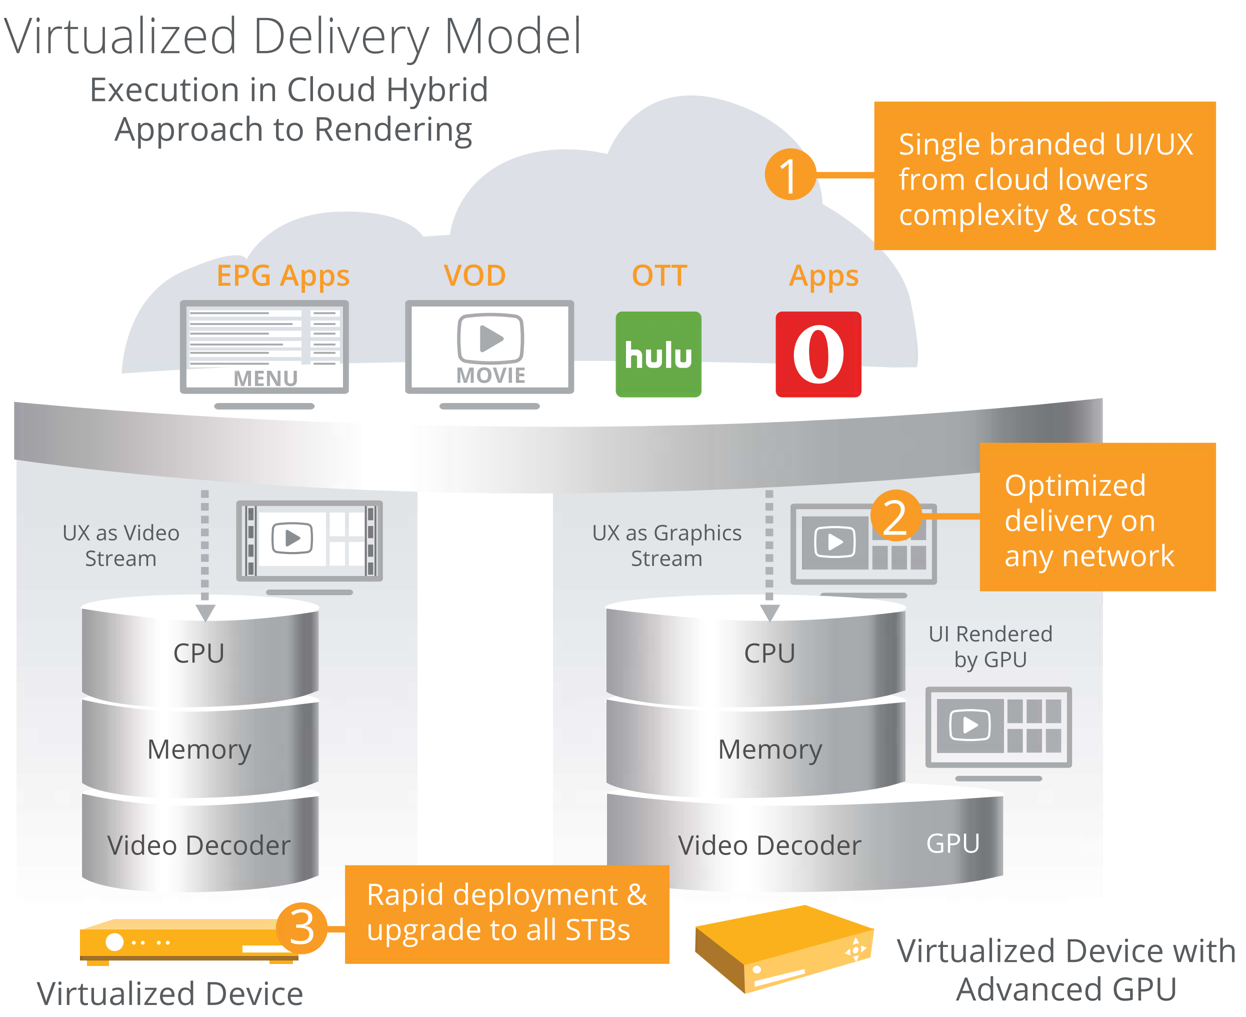

Virtualized Delivery Model

In this image from ActiveVideo,® CloudTV is the core software technology platform that virtualizes device functions by removing the UI/UX dependency from the device and delivers the experiences from the cloud.[3]

The cloud also has much to offer when it comes to advertising, as evidenced by how some Service Providers 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. The result is 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, ensuring 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, Service Providers can send subscribers packaged software that will allow terminals to configure themselves on the network, and to extend configuration to all other devices in accordance with the policies appropriate to each one.

New STB processing requirements

While use of virtualization technology is essential to enabling a next-generation multiscreen TV service, 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 mentioned earlier require new spending; the tasks that remain the responsibility of CPE bring new processing costs into the equation. Even as old tasks are simplified to some degree by the cloud, the best case scenario is a trade-off that leaves the CPE memory and million instructions per second (MIPS) quotients unchanged.

CPE functions 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 high efficiency video coding (HEVC) or other modes of advanced encoding, representing a significant increase in processing requirements over MPEG-2 and MPEG-4 AVC. STBs will also need to be able to support refresh downloads in the event of decoder failures.

Cloud-based personalized UI construction and delivery to each user over IP or MPEG transport feeds can eliminate the need for graphics processing unit (GPU), rendering power in the set-top for both high and standard definition UIs. However, Service Providers are likely to determine that the much higher throughput required for transmitting a fully rendered UHD-caliber UI from the cloud incurs a high bandwidth cost. In this case the primary media gateway in the home would handle the rich graphics rendering required with such UIs. Spending for such capabilities on secondary STBs can be avoided, assuming the media gateway is able to transmit the fully rendered UI over the home network.

Similarly, the removal of the digital video recorder (DVR) from the STB eliminates the traditional personal recording processing and storage requirements. Storage processing supporting IP video caching akin to content delivery network (CDN) type of 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 video on demand (VOD) content.

As STBs are relieved of the cable card burden with the implementation of Downloadable Conditional Access System (DCAS), new requirements tied to securing high-value content, including but not limited to UHD content, will impose new functionalities on STBs. This is reflected in MovieLabs' Enhanced Content Protection (ECP) specifications, which have become the template for all types of premium content as producers take action against mounting levels of piracy worldwide.

Service Providers leveraging state-of-the-art content protection technology are well in line with most of MovieLabs' Secure Media Pipeline (SMP) requirements for end-to-end protection. However, 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."[4]

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.

ENABLING IOT SERVICES

The cloud also has a huge role in IoT development with support for device on-boarding, application provisioning and management, device and user authentication, analytics and much 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, Service Providers can bring a virtually unlimited range of options into the user experience.

Utilizing a wide range of open application programming interfaces (APIs), Service Providers can orchestrate 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 Service Providers can support multiple dynamic business models across complex IoT value chains associated with a wide variety of services and devices. They 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 service offerings.

With all the functionality the cloud can deliver, there still 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 application devices, cameras and sensors and all the data flowing from those devices in accord with the state of each user. Moreover, Service Providers will need to be able 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 Bluetooth Low Energy (BLE).

Clearly, in the IoT domain, as in other areas where CPE virtualization comes into play, memory and CPU requirements will be substantial. Thanks to industry standardization efforts, gateways can now leverage open APIs to support the appropriate type of connectivity for any given device and application, whether it be ZigBee, BLE, Google's IPv6-based thread or Wi-Fi. With such capabilities in place, Service Providers can confidently support a broad choice of end user devices to connect through a single point of communication to the cloud.

REFERENCES

[1] *RDK, RDK Adoption Surges Globally, March 2018*
[2] *CableLabs, Specifications, 2014-2018*
[3] *ActiveVideo, Platform Overview, 2018*
[4] *MovieLabs, Solutions & Specifications, 2013 – 2018*

ABBREVIATIONS

ABR	Adaptive Bitrate
ACS	Auto Configuration Server
AP	Access Point
API	Application Programming Interface
BLE	Bluetooth Low Energy
CDN	Content Delivery Network
CPE	Customer Premises Equipment
CPU	Central Processing Unit
DCAS	Downloadable Conditional Access System
DHCP	Dynamic Host Configuration Protocol
DOCSIS	Data Over Cable Service Interface Specification
DVR	Digital Video Recorder
ECP	Enhanced Content Protection
Gbps	Gigabits Per Second
GHz	Gigahertz
GPU	Graphics Processing Unit
HD	High Definition
HEVC	High Efficiency Video Coding
HTTP	Hypertext Transfer Protocol
IoT	Internet of Things
IP	Internet Protocol
MIPS	Million Instructions Per Second
MoCA®	Multimedia over Coax Alliance
MPEG	Moving Picture Experts Group
MSO	Multiple System Operator
NAT	Network Address Translation
OTT	Over-the-Top
PC	Personal computer
PHY	Physical interface
PON	Passive Optical Network
QAM	Quadrature Amplitude Modulation
QoS	Quality of Service
RDK	Reference Design Kit
RF	Radio frequency
SD	Standard Definition
SMP	Secure Media Pipeline
SNMP	Simple Network Management Protocol
SOAP	Simple Object Access Protocol)
STB	Set-Top Box
UHD	Ultra High Definition
UI	User Interface
VBR	Variable Bitrate
vCPE	Virtualized CPE
VOD	Video on Demand
WiGig	Wireless Gigabit Alliance
XML	Extensible Markup Language