# VIRTUALIZING VIDEO PROCESSING FOR SCALE, AGILITY, AND PERFORMANCE

Glen Griffith, Eric Rosier, Brad Medford
Ericsson, Inc.

*Abstract*

*Video headends are currently built in separate silos for linear broadcast, video on demand (VOD), and multiscreen video services. These service silos, combined with a mix of deployment infrastructures - dedicated hardware, private cloud data centers, and public clouds - as well as a mix of vendors have created a complex and inflexible platform for the processing and distribution of video content. This environment increases the cost and time to market for delivery and monetization of video services.*

*A virtualized video headend based upon open and well-defined APIs provides a platform for expansion and greater monetization of video services and simplifies both the management and operation of video processing across a complex mix of infrastructures and vendors. This type of video processing platform improves service agility by decreasing the time to market for new services, leveraging performance benefits of different deployment technologies, and providing both operational and capital cost benefits.*

*This paper describes the scope and functionality of a virtualized video headend, identifying and detailing the key requirements and the relation to pertinent technologies such as cloud, network functions virtualization (NFV), and software defined networking (SDN). A reference model for a virtualized headend is presented and the associated benefits and challenges are and discussed.*

## INTRODUCTION

The pay TV industry is undergoing a period of dramatic change brought on by a ‘digital revolution’ in video - video content, now available in digital form can be delivered over an increasingly accessible public internet to a growing array of consumer devices, especially mobile devices. Demand for video delivered to mobile devices is expected to grow by around 45 percent annually through to 2020, when it will account for roughly 55 percent of all mobile data traffic [1].

These changes affect all participants in the pay TV value chain, from content creators to broadcasters to pay TV operators to consumers. All players in this value chain are pressured to introduce new video services of higher quality to more consumer devices in a shorter amount of time and for lower costs. Networks for production and delivery of these TV services have traditionally been designed and built as static facilities, sized for peak expected loads, with a fixed topology. Expansion of these networks requires additional capital expenditure, so upgrades are typically infrequent and incremental in nature.

In recent years, the IT industry has undergone a significant evolution towards software and virtualization of resources. NFV and SDN are being incorporated into service provider networks to leverage virtualization and cloud technologies and to create flexible, adaptable network infrastructures for delivery of services. The principles and technologies of NFV and SDN can be extended to video processing in order to bring the same benefits enjoyed by the IT and communications industries. Video processing virtualization – the separation of the management and video processing functions - is applicable to a substantial part of the video processing chain and brings similar benefits as the ones demonstrated in the IT domain.

This paper focuses on the virtualization of video processing functionality for participants in the pay TV value chain, describing the functionality and scope of a virtualized headend, explaining the benefits, describing how it might be constructed, and identifying critical issues that must be addressed.

## TODAY'S VIDEO PROCESSING ENVIRONMENT

Video headends - the portion of pay TV networks which acquire and process video content to be delivered as video services - have evolved in piecemeal fashion. Video headend systems have traditionally been designed and optimized for delivery of linear broadcast television services - the primary type of video services delivered. These systems have been static facilities designed to perform a specific set of functions and sized for a given amount of video traffic. The growth of VOD services in the 2000s saw the introduction of a new type of video processing infrastructure to support the processing of file-based video content. The systems to support this 'offline' processing were deployed as separate video headends, often from a different set of vendors. The rise of multiscreen video services in recent years has created the need for another type of video processing for delivery of video streams to connected devices such as PCs, smartphones, and tablets. This need was filled by deployment of yet another video headend infrastructure, separately provisioned and managed. The traditional pay TV video headend grew dramatically in scope and complexity from the original linear broadcast processing platforms, consisting now of a broad mix of video processing technologies and vendors and requiring as much IT management expertise as video.

Video processing technology has undergone radical changes over the last several years – compression technologies have evolved to the point that multiple standards have been defined and used for a given function. For example, compression technology went through multiple generations of changes - MPEG2 to AVC/H.264 and standard definition (SD) to high definition (HD) for both - in the last 10 years, with the introduction of HEVC on the horizon. Transition to IP-based transport as an alternative to ASI has also brought about extensive architectural changes in the headends themselves and has become widespread, creating additional complexity in the transport formats to be supported.

Another factor driving complexity is the increasing number of transport technologies that must be supported. While MPEG2 transport stream continues to be the predominant protocol, cable, terrestrial, and satellite networks require several different modulation schemes. This, in combination with regional variations, such as the use of Switched Digital Video in US cable, as well as the use of CBR, VBR and ABR techniques have created a complex matrix of use cases.

Additionally, vendors' systems have traditionally been non-interoperable. This results in managing each vendor's equipment separately. The number of Operations & Management (O&M) systems is not only complex but can negatively impact reliability, increasing the Mean Time to Restore (MTTR) system operation.

For these reasons, current generations of video headends have struggled to gracefully cover all use cases and have typically only supported a subset of those required. This is the case in the vast majority of video processing equipment present in networks today.

As a result of these trends, traditional video headends are a mix of different systems built over time to integrate the different advances in video services. The following diagram illustrates today's video headend environment:

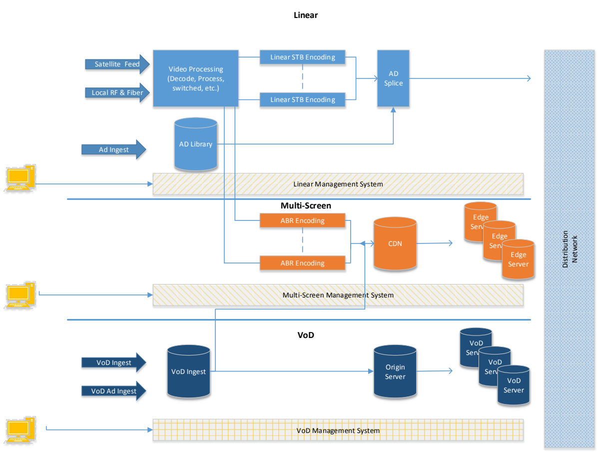


Figure 1: Traditional Video Headend Architecture

The current model for video headends is characterized by multiple silos based on the required services, processing platforms, and delivery networks. Different appliances are needed for different classes of services (linear, video on demand, and multiscreen), as well as for the different deployment technology and transmission networks.

## THE VIRTUALIZED VIDEO HEADEND

A virtualized headend is more than just the implementation of video processing functions in software so that these functions can be virtualized and deployed on datacenter infrastructures. The concept of virtualization for video processing is chiefly related to the separation of the management of video processing operations from the implementation of the video processing functions on a mix of deployment platforms - from purpose-built hardware, Commercial Off The Shelf (COTS) hardware, private cloud, and public cloud infrastructures. A virtualized video headend is capable of managing video processing operations across all available deployment platforms so that each platform can be leveraged for its unique advantages to deliver optimum processing performance in terms of quality, cost and efficiency. This approach creates a flexible pool of resources on which video processing operations may be deployed. A virtualized headend is essentially the application of NFV and SDN to video processing.

The following diagram illustrates the logical layers of a virtualized headend.

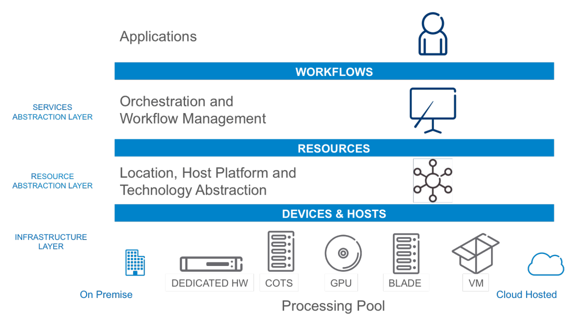


Figure 2: Virtualized Headend Logical Layers

The Infrastructure Layer is where the deployment platforms for the processing components reside. This layer abstracts specifics of the various deployment platforms and publishes its capabilities to the Resource Abstraction Layer. The Infrastructure Layer provides the following:

- Video processing hardware and software
- Video data pipeline - Media Inputs, Media Transformation, Media Outputs
- Low-level Analysis
- Network redundancy ("If this feed is not available, then use another one")

The Resource Abstraction Layer models the processing capabilities, resource requirements, and available resources of the Infrastructure Layer and publishes its capabilities to the Services Abstraction Layer.

This layer provisions and engages appropriate resources to carry out the processing operations directed by the Service Abstraction Layer. The Resource Abstraction Layer provides the following:

- Redundancy of devices and servers
- Resource allocation ("How many of these can I have?")
- License management
- Northbound API for orchestration of processing tasks
- Configuration of system elements, either individual or bulk
- Monitoring of elements ("Is my equipment healthy?")

The Services Abstraction Layer provides an end-to-end service view of video processing and distribution flows from content ingest to delivery, modeling these flows from the perspective of input and output video characteristics and engaging with the Resource Abstraction Layer to effect provisioning and deployment of required processing components on the Infrastructure Layer. The Services Abstraction Layer provides the following capabilities:

- Management interface for defining video flows
- Modeling of video flows based on source and output video characteristics
- Interaction with a Resource Abstraction Layer to effect video processing required by defined video flows
- Workflow management for execution of video processing and distribution rules: **if** this **then** that
- Analytics for management of Service Level Agreements and Key Performance Indicators purposes

## Capabilities of a Virtualized Headend

A virtualized headend must meet many requirements in order to optimize video quality and efficiency of delivered video services while providing the operational flexibility of a virtualized environment:

**Support for all types of video processing deployment platforms**

Video headends consist of a range of deployment platforms. Different types of deployment platforms excel at different types of video processing. Processing of live, broadcast-quality video, for example, requires real-time performance, which is often best supported by purpose-built hardware. Processing of file-based video for a VOD library may be done offline, which lends itself to software processing on COTS or cloud infrastructure where more processing time can be used. Picture quality continues to be the number one rated consumer feature in terms of importance [2], so a virtualized video headend must leverage the advantages of each type of processing platform to deliver the highest quality result for a given resource budget. The following diagram illustrates how a virtualized video headend can harness the strengths of different processing platforms:

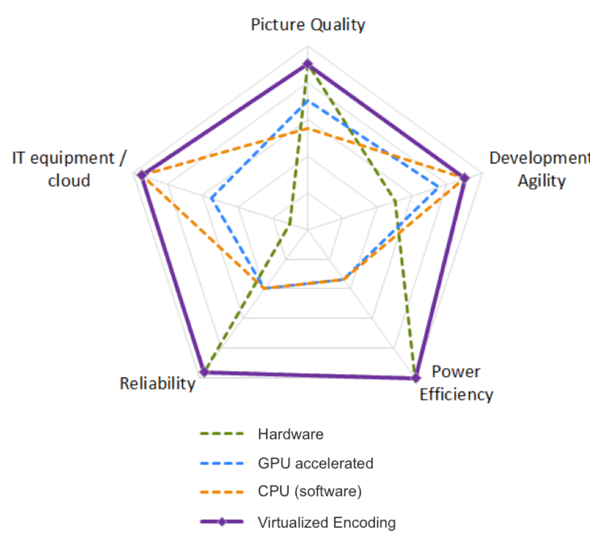


Figure 3: Virtualized Processing Optimization

**Support for broadcast video requirements**

For a broadcaster or a TV service provider, the video headend is the heart of their operation, providing the video content which is the lifeblood of their business. A virtualized video headend must meet the same rigid requirements as traditional broadcast video headends in the areas of availability, content security, low latency, and video encoding quality.

**Converged Management Across Service Types**

Converged management across service types is important, because it reduces complexity and simplifies operation of the video headend. Traditional video headends consist of separate infrastructure silos for linear, on-demand, and multiscreen services, each with its own management system and user interface. A virtualized video headend consolidates the management of video processing for all services into a single platform, creating a service-level view of video processing flows.

**Separation of video processing workflow management from video processing implementation**

This capability is central to increasing service velocity and making the video headend 'future-proof'. Orchestration and management of the video processing workflows must be independent of the processing components to allow changes in processing infrastructure driven by advances in technology and the business/commercial environment.

**Support for multiple licensing schemes**

A virtualized headend introduces new levels of flexibility in video processing workflows. New licensing models are required to take full advantage of this flexibility. At least four licensing modes must be supported:

1) Pay Per feature: Optional base license for an appliance plus licenses per feature, according to different parameters (SD stream, HD, Video Quality, etc.)
2) 'All you can eat': do what you want with the equipment, within the limits of performance of the hardware
3) Per capacity: Use bandwidth (or physical characteristics of transmission) as a dimensioning factor
4) Per use: Pay when you use the system. This is the common licensing scheme of temporary, cloud-based encoding instances

A virtualized headend will require an abstract model of equipment license and unified license management encompassing and aggregating these different models.

**Integration with other network virtualization technologies**

Broadcasters and TV service providers are actively deploying virtualization technologies, such as private cloud infrastructure, software-defined networking, and network functions virtualization to provide dynamic platforms for deployment of network services. Leveraging these technologies maximizes their operational efficiency and allows video processing to build on the strong foundation of these virtualization trends.

**Ability to optimize video processing flows per a processing 'budget'**

A virtualized video headend must model the capabilities and resources of the underlying video processing infrastructure and be capable of engaging those processing resources best suited for the processing operations within a set of constraints or 'budget' for that operation.

**Ability to dynamically allocate logical network 'slices' to meet the processing demands of video service flows**
A virtualized video headend models video processing flows as a collection of virtual functions which can be chained together to create a network slice for the ingest, processing, and delivery of a video service.

**Ability to dynamically scale processing resources with demand**
A virtualized video headend creates an elastic pool of resources from the available infrastructure and is able to dynamically scale processing functions up and down based of the demand from video services.

**Service-Oriented Architecture**
Embracing Service-Oriented Architecture principles - loosely coupled components, each fulfilling its service contract – enables a virtualized video headend to more easily integrate with other systems to extend virtualization and orchestration beyond the video processing domain.

**Open and standards-based**
A virtualized video headend based on open standards enables more rapid development and allows use of innovative new capabilities from the open software community or third party companies. An open architecture, with published APIs provides a platform to develop new capabilities on the virtualized video headend framework.

<u>Virtualized Headend Reference Model</u>

ETSI have defined a reference model for NFV management and orchestration [3]. The following diagram extends this model to include the media processing domain and provides a reference framework for a Virtualized Headend:

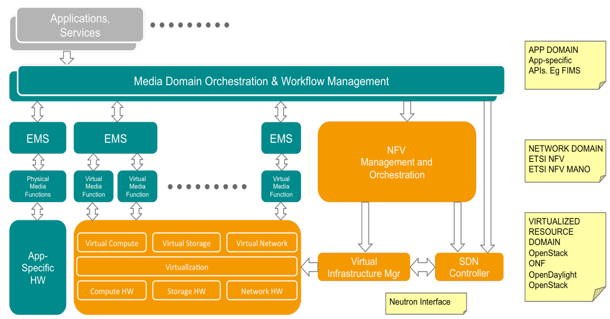


Figure 4: Virtualized Headend Reference Model

Components in green are the components directly involved in providing the processing and management capabilities of the Virtualized Headend; components in orange represent networking and datacenter infrastructure; components in gray are the application and service components which define the video services to be delivered.

Applications and services such as video content management and pay TV middleware define the services to be delivered as well as the parameters and constraints for these services. These service parameters are communicated to the Media Domain Orchestration & Workflow Management layer via a mechanism such as APIs or in-band metadata.

The Media Domain Orchestration & Workflow Management layer distills the service level parameters into required media processing and network functions and orchestrates the instantiation and provisioning of these functions via NFV Management and Orchestration and media Element Management Systems (EMS). Media Domain

Orchestration & Workflow Management may also interact directly with SDN Controllers to request network connectivity for physical functions.

NFV Management and Orchestration configures, coordinates and manages applications, services and underlying virtual and physical infrastructure connected over one or more networks including SDN. Its goal is to ensure the optimum use of virtual resources to meet the required levels of service quality.

NFV Management and Orchestration instantiates required virtual network functions through Virtual Infrastructure Managers and SDN Controllers. VIMs instantiate Virtual Network Functions via a virtualized compute, storage, and network infrastructure.

Media Domain Orchestration & Workflow Management orchestrates provisioning and configuration of the required headend functions on the Virtual Network Function VMs through the EMS components. If physical headend functions are required, Media Domain Orchestration & Workflow Management orchestrates the provisioning and configuration of these functions via the appropriate EMS.

<u>Benefits</u>

A virtualized video headend provides a number of benefits:

**Increased Service Velocity**
Future success for pay TV operators will depend on the ability to introduce new services and features in a timely manner. Consumers have become accustomed to frequent upgrades from app and web-based service providers and expect the same level of service from their pay TV providers. In a virtualized headend the addition of new services or channels becomes more a matter of software configuration than the installation and interconnection of hardware devices. Furthermore, through orchestration and workflow management the software process can be tested and refined and once deployed, can result in new services being rolled out in significantly less time.

**Lower Capital Expense**
As the complexity of video networks has grown so has the associated equipment costs. A virtualized video processing environment lowers capital expense by shifting from a traditional appliance model where processing equipment is purchased then amortized over multiple years to a utility model where video processing resources are consumed on demand. The virtualized headend is also able to manage the video processing infrastructure more efficiently, maximizing the processing output of the available resources thereby lowering overall capital costs.

**Lower Operational Expense**
Traditional video headends require separate management and monitoring systems for different services. The growth in complexity of video processing environments has driven up the operational expense of managing these environments. A virtualized headend is able to decrease this operational expense through converged management of video processing across all delivered services.

**Green**
A virtualized headend provides an environmental benefit through more efficient use of infrastructure, thereby lowering power consumption and reducing required cooling.

**Future-proof**
A virtualized headend protects the investment made in it against obsolescence. Separating the management of the video processing from the

implementation allows processing implementations to evolve independently, allowing advances in processing technology to be introduced while maintaining a consistent management layer.

**Better overall picture quality**
By leveraging the performance benefits of different deployment technologies to exploit the strengths of each, aggregate picture quality across all delivered streams is maximized.

## Challenges

Transitioning from today's traditional video processing architectures to a fully virtualized video headend will present substantial challenges for pay TV operators. Some of these challenges are listed below:

**Legacy Infrastructure**
A flash cut to a completely new video processing platform is not feasible – existing revenue-producing services must be supported; some level of support for legacy processing equipment from a virtualized video processing platform will be required.

**Content Security**
Video content is amongst the most valuable and expensive data an operator will acquire, and content licensing agreements have very strict security requirements. Content security is critical and must be covered according to the expectation of the content owners. As video processing networks become virtualized and employ more COTS and cloud infrastructure, content security must be maintained. Assuming the video processing network is a private network is no longer sufficient to ensure security. Safeguards against unauthorized access must be in place and a content chain of trust must be maintained for virtualized headends to meet the rigid content security requirements.

**System Availability**
Linear broadcast video processing infrastructures are the heart of a pay TV operator's system, with commensurate SLAs and uptime requirements. Traditional IT infrastructures are not built to meet these rigid availability and failover requirements. A virtualized headend must meet these stringent requirements across all types of deployment platforms.

**Dependence on IT technologies**
A virtualized headend places IT technologies in the critical path for all pay TV services – if the IT infrastructure has a fault, it can now affect delivery of video services as well as control and management systems.

**Required skill sets**
A shift to virtualized video processing represents a significant change in the required skills for workers within the video processing environment. Traditionally, workers involved in managing the video headend have had specialized skills specifically related to the video domain. A virtualized headend will require as much IT and datacenter management expertise as video.

## Key Technologies and Standards

### Video Processing Capabilities in General Purpose Processors

Video, driven by the widespread adoption of video-enabled applications and increasing video quality expectations, will dominate data traffic going forward. General purpose CPUs are becoming much more efficient at video processing to cope with their changing workload. Mainstream silicon vendors are addressing this demand. Both Intel and AMD have released several iterations of processor architectures which embed dedicated video encoding and decoding on the processor core. Intel have also released multiple generations of Intel® AVX (Advanced Vector Extensions). AVX is an instruction set extension supported by both Intel and AMD x86 processors which is designed to improve performance of Floating Point-intensive applications, such as image and audio/video processing.

**Encoding Software Tools/Middleware**

Microprocessor manufacturers have enhanced their CPU and GPU architectures to accelerate video encoding and transcoding. In order to maximize performance of these architectures, video processing middleware toolkits such as Open CL and FFMPEG have been developed to provide efficient access to the hardware.

Open Computing Language (OpenCL) is a framework for writing programs that execute across heterogeneous platforms consisting of central processing units (CPUs), graphics processing units (GPUs), digital signal processors (DSPs), field-programmable gate arrays (FPGAs) and other processors. OpenCL includes application programming interfaces (APIs) to control the platform and execute programs on the compute devices. OpenCL is an open standard maintained by the non-profit technology consortium Khronos Group. It has been adopted by Apple, Intel, Qualcomm, Advanced Micro Devices (AMD), Nvidia, Altera, Samsung, Vivante, Imagination Technologies and ARM Holdings.

FFmpeg is an open-source software project that produces libraries and programs for processing multimedia data. FFmpeg includes an audio/video codec library (including HEVC), an audio/video container mux and demux library, and a command line program for transcoding multimedia files. FFmpeg supports a very broad range of audio and video codecs and file formats. Numerous commercial products and services use FFmpeg, including YouTube, Zencoder, and Handbrake as well as many popular game products.

**Mezzanine Format Standardization**

The growth in the use of mezzanine formats for distributing content between broadcasters and TV service providers has led to a need for standardization. Broadcasters do not want to produce a version for each pay TV operator, and pay TV operators do not want to receive a different format from each broadcaster.

Several standardization bodies are currently working in this area:

- The Advanced Media Workflow Association (AMWA) has defined a set of specifications that build on MXF to provide a vendor-neutral exchange format for finished programs.
- In February 2014 a Joint Task Force on File Formats and Media Interoperability was announced by its sponsors, the North American Broadcasters Association (NABA), Advanced Media Workflow Association (AMWA), Society of Motion Picture and Television Engineers (SMPTE), International Association of Broadcast Manufacturers (IABM), American Association of Advertising Agencies (4A's), and Association of National Advertisers (ANA). The European Broadcasting Union (EBU) is participating as an observer.
- The Digital Production Partnership (DPP) is an initiative formed by the UK's public service broadcasters. The group has defined a UK-wide standard for finished program delivery, which is a defined subset of AMWA AS-11.

**Network Functions Virtualization (NFV)**

NFV is an initiative to virtualize the network services that have traditionally been carried out by proprietary, dedicated hardware. The goal of NFV is to decouple network functions from dedicated hardware devices and allow network services that are now being carried out by routers, firewalls, load balancers and other dedicated hardware devices to be hosted on virtual machines (VMs).

**FIMS**

In 2010 the Advanced Media Workflow Association (AMWA) and the European Broadcasting Union (EBU) jointly set up a project to develop standards for a framework to implement a media-friendly Service-Oriented Architecture (SOA) with the aim to provide a flexible and cost-effective, reliable, and future-proof solution for production of broadcast content. The resulting project is the Framework for Interoperable Media Services (FIMS). FIMS describes a vendor-neutral common framework for implementing Interoperable Media Services using a SOA-based system, supporting interoperability, interchangeability and reusability of media-specific services. It should allow best of breed content processing products to be integrated with media business systems.

Example Use Cases

The following serve as examples of how a virtualized headend might operate in practice to address video processing challenges:

**Introduction of an LTE Broadcast service for an existing video content source**

In this scenario, an operator has acquired mobile broadcast rights for a very popular channel. The operator wishes to add this channel to its LTE broadcast service, which requires HEVC encoding. The operator edits the headend configuration for the channel, adding a new LTE broadcast output with desired quality level and processing budget. The virtualized headend provisions an instance of HEVC multiscreen encoding to support the required output profiles and also provisions input and output network flows as well as CDN storage capacity to deliver LTE broadcast service.

**Temporary broadcast of a popular sports tournament**

In this scenario, an operator wishes to establish a set of channels – available in linear SD and HD as well as multiscreen formats – for broadcast of games from a popular sporting tournament taking place over a 3-month period. The operator enters configuration information - video source, desired output formats, and quality levels, as well as the broadcast schedule - into the virtualized headend management interface. The virtualized headend provisions required processing resources from the pool of available resources, selecting the appropriate processing platforms to provide the optimum result. The virtualized headend also interoperates with available SDN controller infrastructure to provision network capacity for ingress and egress video flows.

**Import of a new content provider's VOD movie catalog**

In this scenario, an operator has obtained video on demand movie rights for a major studio's entire film catalog, which includes all titles in both SD and HD as well as streaming rights to connected devices. The operator has launched a marketing campaign promoting the new VOD content. Video from the studio will be delivered in high bitrate mezzanine format. All video content must be ingested, processed, and ready for VOD playout by the target launch date. After the operator configures the requirements of the processing operations, the virtualized headend determines the amount of processing and network resources required to successfully encode the source video files into all desired output formats in the time allowed, storing the content in the appropriate locations for playout.

## CONCLUSION

Advances in network function virtualization and software-defined networking can be applied to the video processing domain to create a virtualized video headend. A virtualized headend abstracts the management of video processing operations from the infrastructure implementation of these functions. Such a virtualized headend must be capable of leveraging multiple types of video processing deployment platforms ranging from purpose-built hardware to COTS hardware to both private and public cloud infrastructures in order to optimize picture quality per a given processing budget. By abstracting the management of the video processing operations from the implementation, the virtualized headend can dynamically provision and allocate processing resources to create video processing 'network slices' in response to the demands of video applications and services. This optimizes overall video processing performance by leveraging the strengths of each type of processing platform and improves cost efficiency by maximizing the use of the available infrastructure and simplifying operations.

## REFERENCES

[1] **Ericsson Mobility Report**
February 2015
http://www.ericsson.com/res/docs/2015/ericsson-mobility-report-feb-2015-interim.pdf

[2] **Ericsson ConsumerLab Report**:
**TV and Media 2014**
http://www.ericsson.com/res/docs/2014/consumerlab/tv-media-2014-ericsson-consumerlab.pdf

[3] **ETSI GS NFV-MAN 001 V1.1.1 (2014-12)**
Network Functions Virtualisation (NFV); Management and Orchestration