TRANSCODING CHOICES FOR 24x7x365 LINEAR VIDEO

A TECHNOLOGY COMPARISON BETWEEN HARDWARE-BASED VIDEO COMPRESSION SYSTEMS AND SOFTWARE-BASED VIRTUALIZED SERVER MODELS

RICHARD PESKE
VICE PRESIDENT OF VIDEO COMPRESSION
AND PROCESSING PRODUCTS

TABLE OF CONTENTS

With the rapid pace of evolution that's taking place in the video delivery marketplace, our industry must constantly evaluate the benefits of new technologies and weight them against those of current methods. This especially holds true for video encoding and transcoding, where traditional hardware-based methods are now being challenged by a new approach that leverages pools of general-purpose servers running software. For many video applications such as VoD, this "virtualized" model is gaining popularity among service providers as an efficient way to accommodate the constantly changing processing loads of those applications. However, the dynamics of linear broadcast video are much different from those of stored content, especially when it comes to transcoding.

This paper examines the unique transcoding requirements of linear broadcast video and compares hardware systems to new virtual server models in supporting these needs. In it, ARRIS examines how each transcoding choice supports the processing workload, channel volume, video quality, and multiscreen delivery requirements of linear video. In addition, we examine the business dynamics of each approach, including the chief drivers of capital and operational costs for each transcoding model. At its conclusion, this paper provides a clear recommendation for service providers seeking guidance on the right technology choice for linear video transcoding.

INTRODUCTION

Consumers have come to rely on network and premium channels as a "go-to" source for entertainment and information, whether they're watching on the living room TV or using a TV Everywhere application on one of the many devices at their fingertips. And with the evolution from standard definition (SD) to high definition (HD) to 4K and UltraHD (UHD), viewers' expectations for a high-quality linear video experience are constantly increasing. For service providers, delivering this experience is a top priority, but doing so must not come at the expense of unreasonable costs and must not exceed the available capacity within their networks. This is where the right transcoding technology can make or break the subscriber experience and the service provider's business case.

With the right transcoding methods, service providers can pack the highest quality video into the fewest bits. This is particularly critical in cases where the bandwidth available for video delivery is a scarce resource. Examples include cable operators that are running out of DOCSIS bandwidth or QAM broadcast spectrum, telecommunications providers using DSL, wireless operators running 4G or even 3G networks, and direct-to-home satellite providers with limited capacity. In addition, many programmers are facing limited transponder bandwidth or seeking new ways to reduce their video transmission costs.

In each of these cases, delivering high quality linear video must not come at the expense of network bandwidth that isn't available or is only attainable at a prohibitive cost. There is also the need to consider operating costs such as power, cooling, and space. Operators are under pressure to keep costs down and to meet new "green" regulations, such as Energy 2020 driven by the SCTE in the US or the energy efficiency directives from the European Commission for 2020 and 2030. And most importantly, the chosen transcoding approach must not negatively impact the subscriber experience. The key to accomplishing all of these lies in the selection of transcoding methods that are tailored perfectly to fit the distinct needs of linear video.

For the purposes of this paper, we will use the term "transcoding" to refer to both the process of encoding uncompressed baseband video signals or transcoding previously compressed signals, unless otherwise stated.

THE UNIQUE TRANSCODING REQUIREMENTS OF 24x7x365 LINEAR VIDEO

Delivering linear video is an always-on endeavor, and that means a steady workload for the transcoders that are preparing each stream for distribution. Much like the news, shopping and sports networks they support, these transcoders are online 24 hours a day, seven days a week, 365 days a year. This differs from the role of transcoding for stored video applications such as Video on Demand (VoD). These applications place a dynamic load on the transcoder, which must process video quickly when, for example, a new movie or television episode is added to the VOD library. Once the transcode is complete, there is a lull in activity until the next piece of content is ready to be processed. When it comes to linear video delivery, there is no rest period.

In addition, transcoders must be able to adapt well to new and updated compression standards. As video distribution has transitioned from SD to HD to UHD, compression algorithms have evolved to deliver this increasing quality without over-taxing the network. With the migration from AVC to HEVC just beginning, transcoders must not only support multiple codecs, but adapt to changes in how the new algorithms are delivered. To support the latest compression standards, transcoders must be powerful enough to handle a significant increase in both the volume and complexity of computations required.

The unique needs of linear video carry with them several requirements for the transcoding methods used to process it. To support the constant needs of 24x7x365 video delivery while maintaining a high quality user experience, today's transcoders must perform at a consistently high level, while being scalable, efficient and adaptable to change. In addition, the right transcoding approach for linear programming must also make economic sense when it comes to capital and operational expenses. Each of these requirements will be used to compare hardware-based systems and software-based virtualized servers for 24x7x365 linear video transcoding.

A TECHNOLOGY COMPARISON OF HARDWARE- AND SOFTWARE-BASED TRANSCODING FOR LINEAR VIDEO

To compare the two most common approaches for linear video transcoding, it is important to first define them.

Hardware-based Transcoding

Broadcast quality video compression has traditionally been performed with hardware systems built with specialized compression silicon. The specialized compression silicon chips used in these systems utilize dedicated circuits to perform the complex and specialized computations of video processing. These systems are built with a variety of specialized interfaces that can ingest video in a range of formats, with the latest transcoder architectures utilizing a modular design that allows these interfaces to be changed as needed.

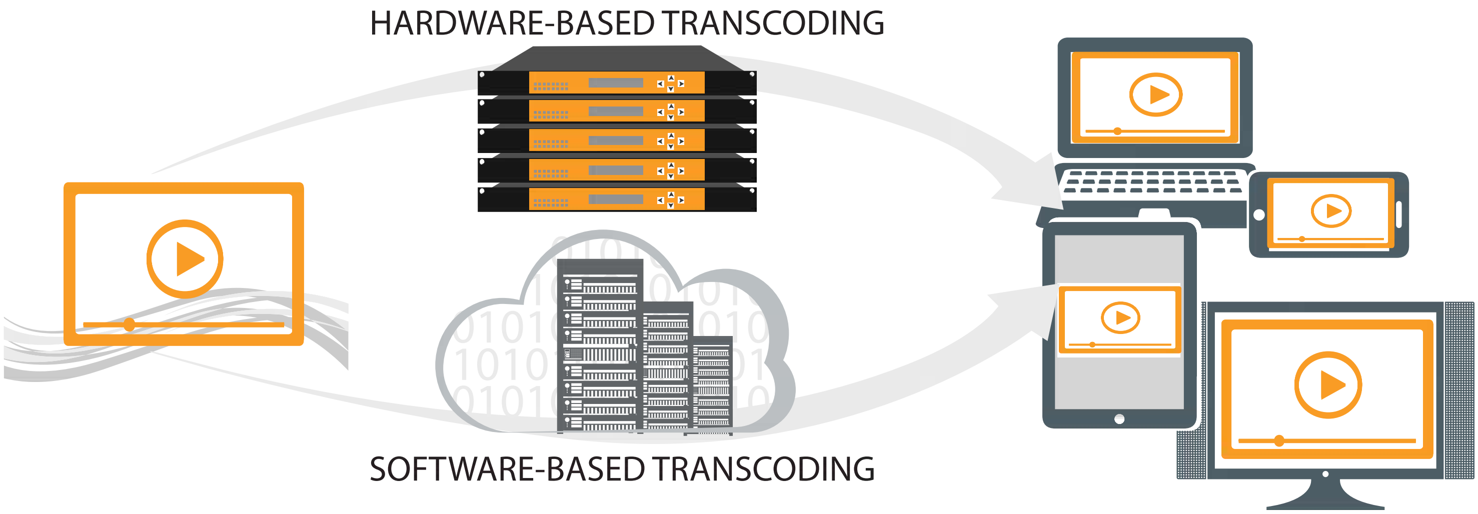

Figure 1: A Simplified View of Hardware-based and Software-based Transcoding Systems

Software-based Transcoding

Software-based transcoding leverages the ever-increasing power of general purpose computing servers to process video. A general purpose server features all-IP interfaces and has a powerful x86 multi-core silicon chip that is designed to handle a wide range of computing tasks. Multiple servers can be connected into a virtualized processing pool, which can be allocated dynamically to the encoding task using software as needed. Server-based transcoding has seen an increase in adoption based on its use of off-the-shelf servers and its flexibility to support rapidly changing workloads.

There is a lot of discussion about using hardware to assist general purpose servers. This can be in the form of added processing cards using either dedicated silicon or Field Programmable Gate Arrays (FPGA). But these types of devices should be classified as hardware-based systems as far as video compression is concerned. Alternatively, Intel is incorporating graphical processing units (GPU) and some fixed function hardware for tasks such as video compression into its CPU processor line. However, servers that utilize these specialized processors are not generally found in standard IT data centers, although their use is expected to increase. In addition, these hardware blocks are typically focused on lower quality Internet-style video, so their use with processing of broadcast quality video for high resolution, large screen display devices is limited. In this paper, we will focus on general purpose datacenter servers used for virtualized cloud applications, unless specifically noted otherwise.

By using the unique linear video transcoding requirements as a guide, we can next analyze how software-based models compare to hardware systems in processing video content for 24x7x365 delivery.

Supporting the Steady-state Workload of “Always-on” Programming

The constant compression requirements of live linear video are quite different from the dynamic workloads typical of the transcoding of stored content or even occasional live events. This need for continuous processing eliminates one of the primary benefits of virtualized transcoding: the efficiency gained from dynamically allocating computing power across multiple virtual machines as processing needs fluctuate. Conceptually, one could adjust the amount of processing power in proportion to the complexity of the video, for example using more CPU resources for complex video sequences, but the overhead of the resource management imposed by virtual machine software is too high for real-time video processing. It is much simpler to allocate a fixed amount of processing resource to a live linear stream, thus negating a major advantage of virtualization.

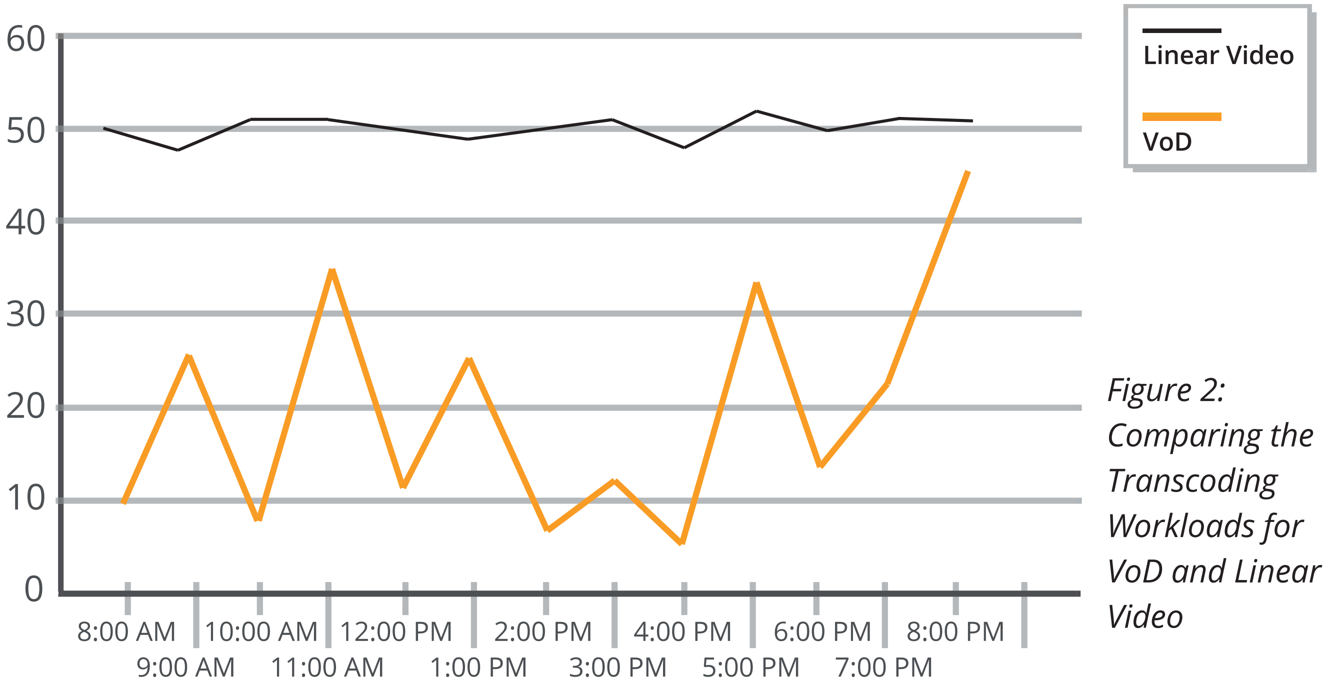

Figure 2: Comparing the Transcoding Workloads for VoD and Linear Video

Conversely, hardware-based transcoding, with its dedicated processing architecture, is much better suited to the steady-state needs of linear video. By maintaining a one-to-one relationship between a linear channel and its transcoding resource, the appropriate amount of computing power can be applied to process video for a given channel. The ability to dedicate resources not only to video encoding but to other processing tasks such as video decoding, audio and metadata, makes hardware-based compression systems much more predictable and makes resource planning much simpler. This predictability is even more important with the addition of the multiple profiles required for multiscreen video delivery. It can be very hard to predict the impact of processing higher bit-rate mezzanine inputs or the addition of a new video profile for adaptive bit-rate (ABR) environments.

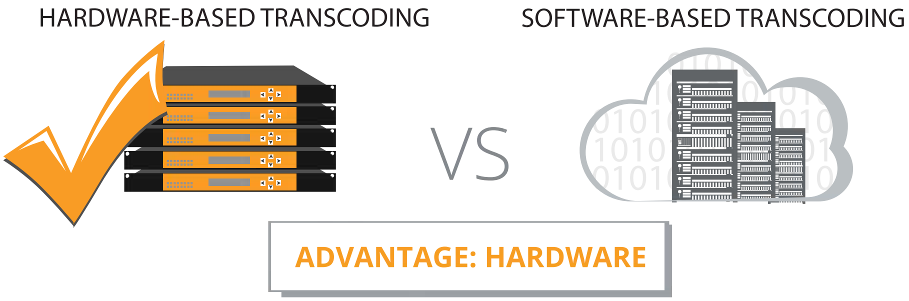

Ingesting Video from Various, Often Geographically Distributed Sources

There are two primary connectivity challenges presented by linear video: collecting video streams efficiently from local programmers or broadcasters that are often a long distance away, and ensuring that they can be ingested easily once they reach the transcoder. In evaluating hardware and software transcoding approaches, these two factors are strongly impacted by the interfaces used to acquire video signals from the network.

Generic servers for virtualized cloud applications use IP interfaces exclusively to ingest video. But delivering uncompressed HD and UHD video streams can require significant amounts of expensive IP network bandwidth. For HD video, each stream requires 1.485 Gbps at 1080i30, and twice that for 1080p60. These requirements increase to 11.88 Gbps for an uncompressed 4Kp60 stream. To deliver high quality video over an IP network cost effectively, these video streams must be pre-compressed. To enable this model, service providers must deploy and manage a series of geographically distributed and dedicated video compression resources, which are not suitable for virtualization in centralized or regional datacenters.

Linear streams are typically encoded and transported from the programmer/broadcaster to the service provider in a high-bitrate mezzanine format in order to maintain the highest possible video quality. Consequently, software-based systems need to dedicate more processing resources to decode these types of streams, which will reduce the resources available

Ifor encoding. Newer hardware-based systems are not typically impacted by the type of decoding required, since the encoding and decoding functions use separate resource pools.

Geographically distributed content and multiple ingest formats are well handled by compression appliances with specialized transcoding hardware platforms, which come equipped with a multitude of inputs to ingest video in any uncompressed, mezzanine or highly compressed format. With a hardware-based transcoder, service providers can customize their transcoding architecture to ingest video in the formats that work best for their programming partners and distribution networks.

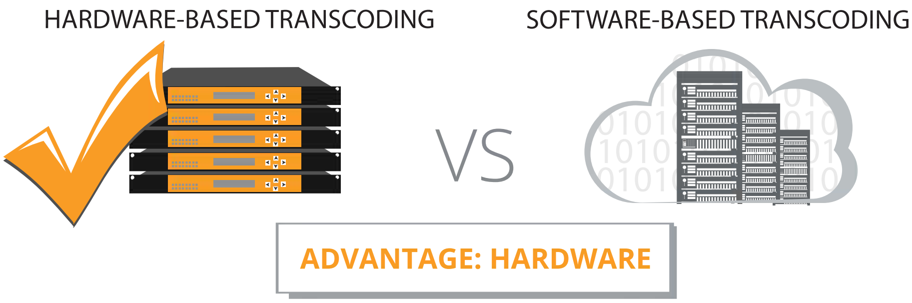

Supporting the Long-term Needs of Linear Video

In the world of video delivery, one thing is certain: linear video isn't disappearing. Live broadcasts will be around as long as there are televised sporting events, 24-hour news networks and pay-per-view specials. But even as a steady force, the delivery of linear video is undergoing change. New compression algorithms are entering the fold with increasing frequency, and existing standards are evolving to be more efficient and stable. When choosing between hardware- and software-based transcoding approaches, we must consider the long-term needs of linear video.

The latest generation hardware-based transcoders are designed and built to last for well over a decade. Fully modular, these systems can evolve without the need for a forklift upgrade or re-cabling in the equipment rack. This commitment to longevity is upheld by silicon suppliers and vendor engineering teams, who understand that video transcoding systems need to remain up-to-date and in production for many years. This extended life cycle is a fundamental differentiator between IT data processing equipment and broadcast video equipment.

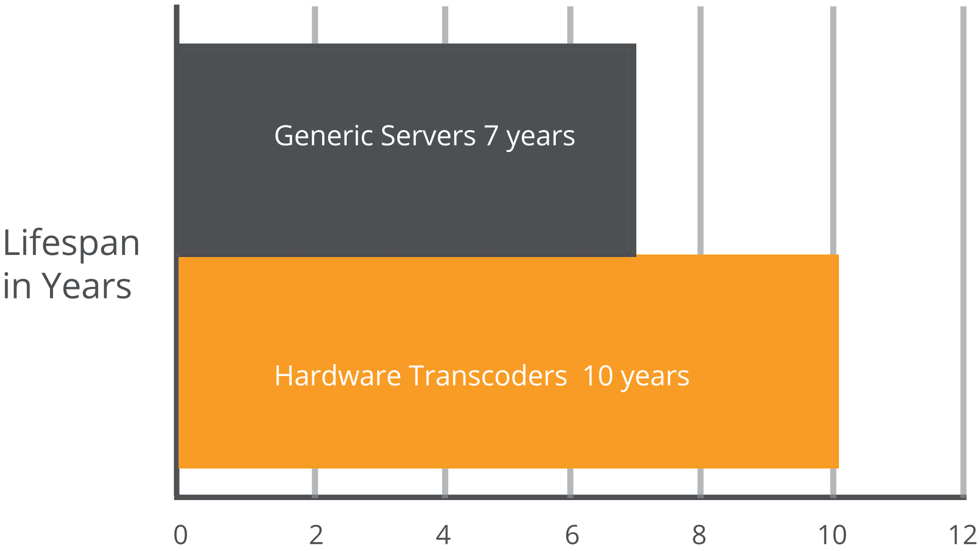

Figure 3: Comparing the Lifespans of Generic Servers and Hardware Transcoders

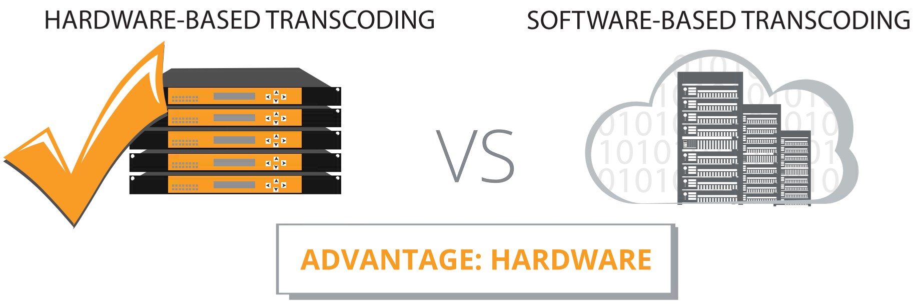

Keeping Capital and Operational Costs Low

In considering the unique transcoding needs of delivering linear video, we cannot exclude the practical requirement for an economically viable solution. And while, on the surface, generic servers can be purchased more inexpensively than specialized transcoding hardware platforms, we must investigate the total costs of ownership for the two models. This calls for an analysis of the drivers for both the capital equipment costs as well as the ongoing operational expenses associated with each transcoding approach.

While generic servers alone can be purchased more inexpensively than hardware-based transcoders, both hardware and software expenses must be factored in for when comparing the capital costs per channel of the total solution. These per-channel solution costs are driven primarily by competitive forces, and therefore are very similar for hardware- and software-based architectures. However, when examining power and real estate utilization for both approaches, the operational expenses of each model are quite distinct.

To compare the operational costs of hardware- and software-based transcoding, we utilize a 5-year (43,800-hour) time horizon, assume a cost per Kilowatt-Hour of $.10 and assign a power utilization efficiency (PUE) of 2.5 for both methods. This metric presumes a slight improvement in PUE for data center operations over the blended industry average of 2.91. In this comparison, we also assume a channel count of 3,500 to demonstrate the compounded operational costs for a service provider with a broad geographical reach and the need to transcode multiple video profiles for multiple display types. We assume the use of AVC/MPEG-4 transcoding as this is the most common format for IP-based delivery.

	2015 TECHNOLOGY	
	Generic Server	**Hardware Transcoder**
Power Dissipation per channel	44	8.3
Cost for one channel	$481.80	$90.89
Cost for 3,500 Channels over 5 Years	$1,686,300.00	$318,097.50
Total Cost Avoidance of Hardware Transcoders over Generic Servers	**$1,368,202.50 saved**	

Table 1: Comparing Operational Costs for Generic Servers and Hardware Transcoders

As Table 1 shows, the power expense savings of hardware transcoding are substantial when compared to generic server implementations, reducing power expenditures by over 80%, a major difference in energy efficiency, which is now a mandate for most operators.

In addition to the power costs of transcoding, real estate utilization is also a critical factor in today's space-limited video delivery facilities. Table 2 compares the rack unit requirements for transcoding the same 3,500 linear channels using generic servers and specialized hardware. Once again, hardware systems provide a significant advantage, achieving a space savings of over 70%.

	2015 TECHNOLOGY	
	Generic Server	**Hardware Transcoder**
Channels per Rack Unit	12	48
Rack Units Required for 3,500 Channels	294	77
Space Savings of Hardware Transcoders over Generic Servers	**217 Rack Units saved**	

Table 2: Comparing Real Estate Utilization for Generic Servers and Hardware Transcoders

The efficiencies of hardware transcoders are clear. In addition to reduced capital costs, hardware implementations consume much less power and take up less floor space than generic servers. While not depicted in this paper, additional opportunities for savings using hardware-based transcoding include reduced cooling expenses, fewer management staff required, and the low maintenance costs associated with the 99.999% reliability of hardware transcoders. And lastly, as the industry begins moving to newer, significantly more computationally complex compression algorithms such as HEVC (typically 5x to 10x the encoding complexity of AVC), the advantages of hardware-based systems are even more magnified.

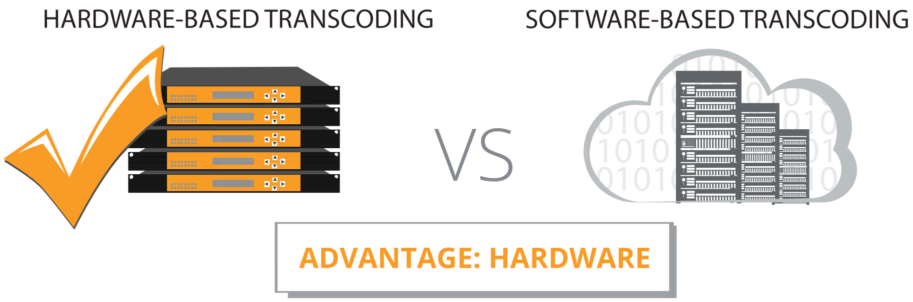

CONCLUSION

When it comes to the unique needs of 24x7x365 linear video, hardware transcoders provide a clear advantage over virtualized software transcoders. Hardware transcoders provide the dedicated processing needed to support the steady-state workloads of "always-on" programming. They are equipped with the input flexibility needed to ingest video from multiple types of sources, while using network bandwidth efficiently. Hardware transcoders also enable the flexibility and longevity needed to keep pace with linear video's ever-evolving needs. And they do it cost-effectively, with lower capital and operational expenses derived from greater density, better power utilization and additional savings associated with greater uptime and decreased cooling requirements when compared to the costs of software-based transcoding.

REFERENCES

1 Niccolai, James. "New data center survey shows mediocre results for energy efficiency." Computerworld, 12 April 2013. Web 24 Mar. 2014.

MEET OUR EXPERT: RICHARD PESKE

As Vice President of Video Compression and Processing Products at ARRIS, Richard leads the product management team for the company's multiscreen and next-generation HEVC solutions. In this role, he drives the definition and development of advanced hardware and software products that maximize video quality for end users, while minimizing the costs, complexities and network impacts of video delivery for service providers. Richard's 20+ years of experience in digital video delivery includes industry pioneers such as BigBand Networks, DiviCom, Silicon Graphics and Bell Labs. He holds an MBA from the Kellogg School of Management at Northwestern University, an M.S degree in Electrical Engineering from the Illinois Institute of Technology and a B.S degree in Electrical Engineering from the University of Michigan