TECHN SITE

A Technical magazine of CSE Department

Volume 5 Issue 2 May 2016

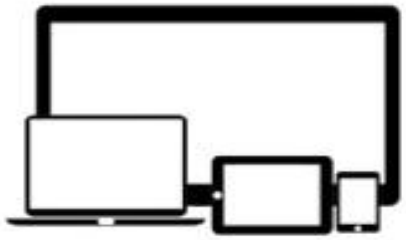

Technical Articles

Article-1

Cloud Computing

Migration to cloud Across the entire content delivery industry, development of new business models is being driven by changes in consumer viewing habits and by the emergence of new technologies, particularly powerful cloud-based solutions that address content source selection, service delivery and user interfaces.

The combination of increased data speeds, new digital compression techniques, increased accessibility of content and alternative consumer premise equipment (CPE) develop into nonlinear services with smart functions such as Search and Recommendation, and cloud-based TV Everywhere services enabling subscribers to view any content on any device at anytime, and anywhere.

Pay-TV operators, who often also provide the internet infrastructure to the home, are uniquely positioned to deliver the emerging aggregate viewing experience. In adopting the cloud model, the operator's focus shifts from installation and maintenance of a large hardware infrastructure to management of virtual resources through software paradigms. This shift offers advantages in cost savings, operational efficiencies as well as improved customer products and services.

The cloud model is based on virtualization, the ability of software to describe functionality in abstract terms, independently of the server platform on which it is hosted. Cloud-based technologies can address not only the current marketplace, but also provide ongoing elasticity for accommodating additional changes that will inevitably emerge through sources of video content, means of delivery and consumers' viewing preferences.

Benefits of Adopting the Cloud Model

In a traditional media delivery environment the hardware required for the content and management infrastructure is procured and managed by the operator on the operator's premises. This scenario requires on-site expertise for upgrading hardware, managing failures and performing fairly complicated software and hardware upgrade procedures.

The cost of hardware, electricity and environmental support are all ongoing expenses for the operator. The cloud computing model provides the cloud with active abilities, with storage and processing power which can scale with agility, both upward and downward, in accordance with demand. Cloud-empowered providers can draw on as much or as little of those resources as needed on a moment-by-moment basis. In the cloud model, the operator has many options for ensuring sufficient resources, both on the server side and on the client side. Furthermore, network support can be virtualized in the cloud, allowing for load balancing and other functions that take advantage of the cloud's flexibility.

Client-Side Changes

The consumer's 'aggregated' viewing habits draw upon video content not only from real-time broadcasts but from such sources as VOD services, OTT and IPTV, gaming, and user generated content (UGC) libraries. In some households, the means of delivering the aggregate client experience has shifted from a single purpose dedicated set-top box to an in-home gateway, providing multiple UIs for delivery on TV, tablet and smartphones, presenting content streamed from the traditional cable feed and the Internet.

Virtualized UI delivery based in the cloud can improve operations by providing a thin client delivery model, extending the life of legacy STBs and enabling new lower cost platforms. This virtualized server-client UI takes the burden off the STB, provides more feature-rich UIs and opens the opportunity for different UIs on different devices. For example, a child's device (whether an STB in a room or a tablet), can use a specific child-friendly, parental controlled UI. Virtualized cloud-based UI delivery also enables operators to target specific groups of STBs for trials or to send down diagnostic UIs for testing a box, reducing truck rolls.

When the service provider places as much intelligence as possible in the cloud, not only is the process of developing and testing and deploying

UIs simpler, and thus less expensive, but further cost savings are achieved by reducing the requirements for the STB.

Server-Side Changes

Technological advances in virtual server infrastructure development are rapidly creating more flexible, dynamic and reliable operations. Virtual server architectures provide the flexibility to address all the needs of the operator, from content delivery to user management, to logging and system monitoring and management. Each specific task area can expand or contract as needed.

A virtual architecture puts the cost and burden of upkeep on the data service provider; because these virtual cloud computing centers are designed with redundant systems and dynamic expansion capabilities, they remove the operator's burden of projecting additional servers or upgrades to their own systems. This allows an operator to keep some operations local, using the cloud infrastructure for others, and answers the need to maintain disparate server locations for failover.

Shifting storage to the cloud reduces the need to maintain local disk storage and its accompanying backup and redundant systems. Reliability, failover and load-balancing functions are handled virtually and in some cases, by applications, particularly in "platform as a service" environments. Proprietary software costs are reduced thanks to the availability of open software solutions for monitoring and managing services. Software updates are easily propagated over the cloud managed network making it easier to keep software up-to-date.

Operations become more efficient by moving back office operations to the cloud. The MSO no longer needs as many employees to physically rack and connect equipment, maintain air conditioning, etc. Of course, while much reduced, the cost of operations does not entirely disappear. Capital expense for new servers may be replaced by cloud storage provider expenses. There may be fewer traditional IT staff, and there will certainly be a higher ratio of servers to staff. Applications running in the cloud must maintain high standards of reliability and fault tolerance.

These functions, once the domain of hardware, in the virtualized cloud environment become the responsibility of software, thus some IT functions may shift to the cloud provider, and some may be assumed by applications.

Network & Delivery Changes

Streaming of on-demand content is growing at a steady pace, and half of viewers routinely use their computers and internet connections for video delivery. The cloud model improves the efficiency of delivering the service to any device in any location (TV everywhere). For example, Cloud DVR allows subscribers to schedule recordings, giving them the option to watch video anywhere, anytime, on a variety of devices. Users

can store other personal video in the cloud for easy access and MSOs can charge for storage, adding a new revenue stream for their business.

Software-centric video service delivery systems provide an advantage that propels operators beyond the current limitations of their physical resources and enables them to better respond to new standards, new opportunities and changing market dynamics.

SVNNS Sai Madhava
(12K61A0596)

Aritcle-2

The challenges of 4k and UHDTV1

The implementation of 4k and the UHDTV1 4k profile as defined by SMPTE ST 2036- 1 could be regarded as an evolutionary process, given that the majority of today's implementations are deployed using multi-link 3G SDI. However, it must be considered that the emergence of high frame rate (HFR) 4k/UHD as well as high dynamic range (HDR) will at least double data rates, demanding completely new video infrastructure.

Many consider that given this, the adoption of IP infrastructure will be the logical approach for accommodating HFR 4K/UHD production workflows, rather than extending the life of SDI with higher data rate implementations such as multi-link 12G. All-IP live production workflows are regarded by many as a revolutionary change that will demand entirely new skillsets and infrastructure. Such a revolution can also enable fundamental changes to the production workflow that will in turn support new business models and services.

It is clear that these changes will impact all aspects of the broadcast chain including content production, broadcasting, content providers, video service providers and equipment (including test and measurement) manufacturers. As stated previously, today's 4k/ UHDTV1 can be accommodated in live production environments with existing 3G infrastructures by using a dual or quad link arrangement (depending on frame rate). However, to deliver 4k/UHD to viewers at acceptable bitrates, ideally requires the adoption of a more efficient codec than H.264/AVC. This is leading to the adoption of the High Efficiency Video Codec (HEVC) otherwise known as H.265 which can offer up to 50% bitrate savings versus H.264 for equivalent video quality.

When it comes to testing and monitoring 4k/ UHD infrastructure the first challenge for test equipment vendors is to support the new standards, ideally with existing tools where possible. Although the availability of 4k sensors and display panels has accelerated the drive towards 4k/UHD implementation, 4k has been around, at least in demonstration form, for almost ten years. Indeed when developing the WFM & WVR8000 series waveform monitors the internal architecture was implemented to handle 4k data rates, which allowed us to deliver 4k/UHD support as an existing product upgrade, rather than requiring a completely new product.

Likewise encoder, transcoder and decoder developers as well as operators need analysis tools to enable development and evaluation of new products to support HEVC/H.265. The architectures of both our H.264/AVC codec analyser and our picture quality analyser were both designed from the outset to handle future codec and format upgrades as they became standardised. This enabled Tektronix to implement both HEVC/H.265 and 4K support as existing product upgrades, rather than as new dedicated analyser products.

Test challenges with 4k/UHD workflows

When deploying 4k/UHD, let us first consider production, post production and broadcast engineering. In these applications, the main test challenges with 4k/UHD workflows today are twofold: the wide variety of formats that need to be supported; and how the timing relationships are maintained between the multiple SDI links. For 3840x2160 or 4096x2160 resolutions at frame rates from 24 to 30Hz these can be carried as Quad link HD or Dual link 3G.

However these low frame rates can produce noticeable temporal artefacts. Therefore there is a need to produce content in 48Hz, 50Hz and 60Hz, which require quad-link 3G SDI arrangements. It is anticipated that in the future, 100Hz and 120Hz frames rates will be adopted which will require support for 24 Gbps data rates, thus making the use of existing infrastructure non-viable.

At such high data rates, the new infrastructure would need to be implemented as dual link 12G SDI or 40G IP. It is therefore quite possible that with the adoption of high frame rate 4k and UHD, we will see the eventual replacement of uncompressed baseband with light compression techniques that would allow the use of readily available 10G IP infrastructure.

This however is in the future and at this time the need is to keep the four SDI signals time aligned. This is critical as each SDI signal flows along a different routing path. Care has to be taken to ensure the timing at each input to a device is within limits and the user has to understand what synchronisation can be accommodated by the device internally to null out timing errors.

In addition to validating correct timing it is important to check the Video Payload ID information to validate that the format is correct and that all links are correctly connected to the device. If present, byte 4 shows the multi-link information for the links and allows the user to verify if the links are in the correct order. 4k and UHD also support an extended colour gamut with the use of ITU-R BT. 2020.

This extends the colour gamut by around 30% when compared to the BT. 709 HD colour space. It is possible to use the data bits in the SMPTE 352 payload ID to automatically detect the colour space or indeed to force the colour space between BT. 709 and BT. 2020. This will slightly change the waveform monitor's vector and lightning displays and also the conversion between YPbPr and RGB on the waveform display.

R Naga Lakshmi Devi
(12K61A0588)

Article-3

Multiple Screens

Over the course of the last decade, we have seen an explosion in the number of devices on which we can view video content, with the average UK household now reportedly having six or more connected devices. There are now thousands of connected devices running on a range of different platforms that have changed forever the way we consume and interact with content. This proliferation of devices has made it difficult for operators to provide multiscreen services that have the capability to adapt to and work perfectly on all devices. In order for any operator in this space to succeed, they must provide multiscreen services that support a wide range of browsers, operating systems, screen sizes and security specifications.

This support is essential as operators deploy multiscreen solutions to complement their main TV offering. These complementary TV services have shifted from an additional extra to a consumer expectation and the operators who meet this multiscreen challenge best will be the ones benefitting from increased loyalty and reduced churn. Operators must re-evaluate their security infrastructure in this new multiscreen space and ensure that it meets the studios' requirements for the avoidance of unauthorised streaming of content alongside the protection of personal data that allows consumers to access services from multiple devices of their choice.

Multiscreen services are causing operators to rethink their delivery methods and to choose how to deal with this new mass delivery of TV Everywhere services, which how includes delivering true TV Everywhere in and out of the home. There are now a variety of distribution networks to choose from such as in-home, cloud, Wi-Fi and mobile and this increased selection is causing operators to amend their infrastructure.

Cloud TV services have proven to be a popular choice as operators seek the best ways to deploy their multiscreen offerings. However, there are question marks over the reliability of cloud TV services and the Quality of Service (QoS) they provide. As more and more consumers use these

services, there is an increased need to ensure that buffering or drops in quality are kept to a minimum or eliminated altogether, especially from providers offering premium services including Hollywood movies or live sports.

Solving the fragmentation issue

The future of TV will encompass traditional broadcast services, IP delivered content and access to personal content, whether recorded on one or more in-home Personal Video Recorders (PVRs), in the cloud via a PVR, stored on personal Network Attached Storage (NAS) or on tablets and smartphones. In order for this new TV landscape to flourish, the fragmentation issue must be tackled head on as operating systems, screen sizes, storage capacities and network access capabilities vary on each device.

Areas in which operators face fragmentation issues when implementing multiscreen include:

• Security and Rights negotiation;
• Device type;
• Device OS;
• Future browser fragmentation;
• Screen size;
• Interaction method (remote, keyboard, mouse, touchscreen, etc);
• Payment / billing / reporting / etc; and
• Advertising support.

These fragmentation issues hinder the implementation of multiscreen and these must be addressed if they are to ensure a seamless experience on all devices. Standards play a key role in enabling devices to be interoperable and seamlessly support consistent consumer experiences that enable operators to increase loyalty, reduce customer churn and win new customers.

One essential element of providing universal solutions that work across all devices is through support for diverse Conditional Access (CA) and Digital Rights Management (DRM) alongside flexible UI technologies and more efficient use of the IP backbone. The unification of CA and DRM is especially crucial for operators in providing a robust multiscreen security solution. Standards such as HTML5 coupled with responsive design enable operators to design a user experience (UX) that can dynamically adapt to the screen size of the consumer's device of choice while new extensions including encrypted media extensions (EME) and the media source extensions (MSE) ensure that access to content is secure. This enables operators to safely deliver high value content direct to the browser.

The multiscreen phenomenon is reshaping the TV landscape with many new distribution networks utilised in order to transmit this content. The new DLNA Commercial Video Profile 2 (CVP-2) guidelines also have an important role to play in ensuring the safe delivery of content across these multiple networks. They offer a fully secure and interoperable solution for the in-home distribution of content.

The guidelines allow for A device discovery and remote user interfaces (RUIs), easing consumer adoption and promoting the operator's brand to unmanaged devices without the need to write a new app for every screen. CVP-2, when combined with DTCP-IP link protection and CA/DRM systems enables operators to deliver high quality content to more devices throughout the home. All these standards work together to allow operators to encrypt content across the open network though there is an onus on them to ensure that they "trans-crypt" from one protection solution to another to avoid compromising on security when consumers share content within the home.

Advanced multiscreen solutions leveraging standards enable operators to monetize multiscreen through advertising insertion and data collection capabilities. These innovative solutions can also provide video metadata content aggregation that can help operators provide additional content tailored to the end user's viewing habits. This ability to better understand the consumer through data analytics such as personalised recommendations and targeted advertising has been proven to increase customer satisfaction and revenue or ARPU (average revenue per user). The integration of standards based technology means that operators can pave the way for new advertising revenues through leveraging new branding opportunities, including interactive banners or video advertisement in the first impression screen or as pre- or post-roll video.

To be successful in the multiscreen space, it is essential that operators, broadcasters or other companies provide a service that encompasses access to as much content as possible, and across as many devices as possible. Standards-based software will be key in making this a reality through its ability to provide operators with the dynamic, flexible and secure solutions that the future of multiscreen demands.

B Harshavardhan Reddy
(12K61A0514)

Article-4

Cloud based DRM

As the video delivery service industry as a whole becomes more software-based, the implementation of cloudcentric security management would seem a logical progression. However, the issue is complex and requires some unpicking.

In some ways, the term cloud is a misnomer. Effectively, DRM delivered as a service from a scalable platform allows content owners and providers to shift away from the need to build, maintain and ultimately upgrade the DRM platforms that are essential for the majority of service provision. “In the past, providers would build their own DRM platforms which would effectively become a software platform needing maintenance and upgrades in line with changes to the underpinning technologies and business evolution,” explains Giorgio Tornielli, VP product engineering, Piksel.

“The cloud model shifts to a case of integration and customisation with the ongoing platform development and expansion taken care of by the service provider. This approach benefits from economies of scale, specialisation and the ability to provide more ancillary features that would be prohibitively expensive to develop for a single in-house installation.” Several consumer DRM products (PlayReady, Adobe or Widevine for instance) are already offered as a cloud service. Cloud DRM is also an emerging trend for implementing operator DRM products, and, in particular for smaller operators, suggests Nagra’s senior product marketing director, Christopher Schouten, as cloud services are an efficient way to optimise their own operations while getting a robust level of service.

Looking further, Schouten sees large operators “starting to run their own cloud infrastructure, implementing a generic IT infrastructure server virtualization approach that brings cost benefits as well as added elasticity.” Such an approach, he informs, “usually relies on a combination of a private cloud environment with some highly scalable online services put on a public cloud or using third party cloud services, for functions such as multi-DRM management.”

Just as functionality like middleware and user interface are moving to the cloud from the client, so security functionality is also ready to make that transition, according to Leonid Sandler, CTO DRM, Cisco. “This is a trend in the beginning stages with growth potential,” he says. There is a broad reason why migration of DRM in hosted or cloud environments for pay-TV operators is occurring: namely, there are core benefits shared with other cloud technologies. As discussed, it is part of a wholesale shift of the video prep and delivery infrastructure towards pure software subsystems, fundamentally dependent on IP connectivity, which can then be implemented in physical data centres or in cloud CPU resources.

In this case, according to Steve Christian, SVP Marketing, Verimatrix, “DRM technology and business logic does not change.” Some vendors have taken generic implementations of DRM technologies (Verimatrix is one) and are offering a SaaS model of key management based on cloud resources. This, says Christian, “tends to shift the business model for DRM costs rather than the underlying technology.”

Cloud-based DRM is an effort to address subscriber interest in viewing online content that is encrypted in various DRM formats and is made available through multiple online video providers. This results in operators needing to support different content packaging and content protection formats like MPEG-DASH and Common Encryption Scheme (CENC).

As more device functionality is moving to the cloud (user interfaces, DVR storage, preference management etc) vendors see a greater share of the core rights management logic moving upstream. "Traditional pay-TV operators have been limited to supporting a single DRM vendor given the operational complexity in rolling out multi-DRM library support into their device footprint," says Sachin Sathaye, VP, product strategy and marketing, ActiveVideo. "The compliance rules associated with supporting DRM vendors further inhibit expanding the DRM solutions to existing devices that lack sophisticated cryptographic features." Moreover, according to Cisco's Sandler, the main advantage of adopting cloud DRM is the moving of license keys and many security functions from the client device to the server.

This allows easier management, updates, and support for multiple DRMs, formats, protocols and evolving business models. Content owners will continue to encrypt and package the same way, however, with licenses and metadata information passed to the cloud DRM server for authentication. They will also be able to offer an extensible service, across regions with availability of servers closer to the end customers.

Security and business drawbacks

Deploying DRM in the cloud in a manner that is secure and in agreement with the DRM vendors' studio compliance rules, is not a trivial task, and cloud security is still not seen as robust as data centre security. Key security functionality such as DRM requires a robust, trusted and scalable provider that has a history of studio approvals and uncompromised security delivery. "It is very important that DRM infrastructure deployed in the cloud is done so in a manner that ensures studio compliance," stresses Steve Plunkett, CTO, Red Bee Media.

"The storage of key material is subject to particular restrictions that must be met to ensure compliance. It also makes more sense to use cloud based DRM when media processing workloads in general are moving into the cloud. Doing so in isolation, while general media processing remains on premise, can slow down the media transit path and increase workflow execution times."

The failure to provide a seamless user experience when wanting to transfer viewing midstream of Netflix' House of Cards from an iPad to a Samsung smartTV, at best risks alienating the subscriber to another service, at worse pushes them toward pirate sites. "The real business challenge for operators who wish to target the broadest number of devices is

to eliminate service distribution and consumption silos that serve only to frustrate consumers and may nudge them towards alternative sources," says Christian. "One aspect of this is the need to enable support for multiple native DRM systems on the devices and browsers in use, and to provide the user with a transparent consumption experience.

"The purpose of security goes far beyond the defensive aspects of addressing piracy and theft of service, says Christian, "to that of ultimately enhancing the subscriber's quality of experience (QoE) while also underpinning the operator's bottom line." Partly for this reason we've not yet seen the wholesale movement of DRM into virtualised environments. "The cloud aspect of DRM is not a process that aggregates large amounts of sensitive customer information, billing or account accessibility," explains Piksel's Tornielli.

"In essence, it is a streamlining of a well understood process that has typically resided inhouse and can now be consumed as an externalised service due to the maturity of connectivity and virtualised computing technologies." André Roy, head of security practice, Farncombe warns that not all cloud based DRMs are optimised to handle content preparation. "We audit CA systems and can say that there are few cloudbased DRM systems that meet studio content handling requirements for handling unencrypted content. Primarily, having studio quality mezzanine files unencrypted in the cloud would not meet MPAA (Motion Picture Association of America) requirements.

All encryption would need handling in a physical facility rather than in the cloud." Verimatrix suggest that full DRM systems (generally consisting of digital rights to manage, encryption, license management, and a DRMenabled client) can be deployed in the cloud "very easily as a part of an overall video delivery system" but it doesn't see this as the gold standard right now for payTV.

"Physical data centre deployment and virtual machine deployments are still quite popular, although interest in cloud implementations is growing," reports Christian. Another potential drawback is flexibility. In certain cases, cloud based services will only support a limited set of DRM technologies. "If a particular feature is not available then it is not always possible to just build it yourself and add it to the mix," reports Tornielli. "The cloud can be defined as a shared multitenant environment so the timescale of upgrades, features and fixes are dependent on the diligence of the provider."

Piksel suggests cloud DRM solutions are a relatively small percentage of deployments (probably less than 1%) "although growing faster than on premise implementations." "Service providers may reduce costs by outsourcing, but also lose considerable control of a key functionality like content security authentication and policy enforcement,"

warns Cisco's Sandler. He adds that cloud security is still not as robust as data centre security, which requires cloud components and deployments to be designed to minimise the associated risks. Only a partial shift of the DRM infrastructure to the cloud is expected from Telekom Innovation Laboratories, the research wing of Deutsche Telekom.

"This shiftimplies the adaptation of the current DRM infrastructure towards more centrally hosted key servers and therefore an overall reduction in the number of key server APIs for the service operator," says Dr Oliver Friedrich. "It also implies the adaptation of content servers hosting files encrypted by means of common encryption scheme, which could also simplify the DRM encryption infrastructure. "Moreover, support for interoperability is currently not being offered by most market leading DRM providers." Therefore, he says, a real move into the cloud is not taking place. The primary DRM innovation for Friedrich, is driven by the browser and multidevice scenarios and technologies, such as MPEG DASH with CENC and HTML5 EME. The most important factor is the decoupling of the content from the DRM system itself.

For an organisation with no existing investment such as a new OTT, SVoD or TVoD entrant, it's hard not to make the case for cloud based DRM from day one.

K. Narayana Rao
(13K65A0501)