Abstract

Linear broadcast has been under constant pressure from new forms of television such as OTT, video-on-demand and streaming services that have driven significant and continuing changes in the way consumers discover, access and view content. Cloud infrastructure and Internet technologies have been at the heart of these disruptive new entrants because of the unprecedented scalability and reach they offer. To stay competitive and augment, if not replace, their linear heritage, it's time for broadcasters to embrace cloud technologies to secure their place in the future of television by becoming more flexible in their ability to deliver high quality, curated content to increasingly fragmented audiences. Playout solutions developed specifically for the cloud now provide flexible tools that are credible, reliable replacements for traditional hardware-based technologies.

Summary

In this paper we aim to demystify key aspects of cloud computing as they apply to linear playout and related linear television technologies. With rapid changes in video content monetisation and distribution business models, it is inevitable that the broadcast industry will have to adopt cloud technologies and resources for playout and the rest of their workflow.

To maximise the value-added proposition of the cloud, decision makers need to understand the key benefits any true cloud solution should bring – accessibility, elasticity, redundancy and resource pooling capabilities.

There are multiple benefits to using cloud playout instead of traditional, hardware or software-based playout systems:

1. **Automatically Redundant Playout:** Redundancy and resiliency are direct, in-built by-products of cloud-based solutions;
2. **Performance and Scaling is a non-issue:** Performance can be increased with no real limitations by scaling the number of services and servers, or the actual performance of each separate server;
3. **Easy Deployment and Use:** In the case of SaaS systems, a broadcaster never has to think about the underlying infrastructure that runs his playout services.
4. **Global Access:** All of the services are accessible from anywhere, usually through a simple web browser interface.
5. **Lower Costs and No Upfront Investment:**
    - Setting up broadcast services from the cloud requires no initial investment.
    - Total *ongoing costs* (including a broad range of costs related to the operation of hardware in house or "in the system") are generally much lower than in traditional broadcast.
    - Overall cloud resource costs are going down as a result of competition and consolidation between cloud resource providers.

Playout History and Challenges

For decades, broadcast engineers and linear playout technology providers focused primarily on sustaining or improving high standards of reliability and stability, which has been achieved by building complex and redundant hardware/SDI/ASI-based broadcast infrastructures.

High capital expenditure costs and the complexity of implementing playout workflows has been one of the consistent factors in assessing the financial viability of the linear television business for broadcasters and their playout service providers. With the industry focus and revenue growth now shifting from linear to digital media/VOD delivery platforms, large upfront investments in traditional linear television playout systems have understandably been questioned, and often deferred, by senior executives and strategists. As a result, broadcasters and playout service providers are often saddled with dated and complex multiple-supplier technologies in their workflow with limited capacity to justify substantial upgrades to deal with new delivery expectations. Forward-looking business models cannot support high spends on redundant and complex internal solutions, and now favour standardised and elastic cloud-based technology.

In the past 10-15 years, linear broadcast technology has gradually evolved from traditional playout hardware solutions, which still dominate the space, to more affordable, but often limited in functionality channel-in-a-box systems. In recent years the trend of virtualising playout applications by moving them from proprietary hardware to generic server platforms has gained momentum as broadcasters look to it as a means of reducing infrastructure and operating costs. Overall, the installed base of playout technologies deliver functionality and reliability that a broadcaster has historically required, but operational flexibility and speed are proving to be the keys for the long-term competitiveness of linear television and several aspects of traditional technologies constrain the ability to expand the flexibility and attain the speeds required to address new market realities. The top two constraints are usually:

1. Heavy reliance on proprietary SDI and GPU video cards;
2. Lack of sufficient IP signal support in head-end infrastructures;

These technologies no longer deliver future-proof solutions. Most other mission-critical IT solutions have steadily moved from hardware to software / virtualisation / cloud based solutions, but playout systems have not, largely because they remain so hardware dependent.

This is a problem because the explosion of Internet delivered video-on-demand content, the fragmentation of audiences, multiscreen consumption and the emergence of OTT as new distribution channels have created major challenges that hardware-based playout solutions cannot resolve.

This transition creates major challenges for the linear television industry because it requires a substantial rethink in infrastructure and associated technologies with new, more flexible broadcast business models. Existing software running traditional or channel-in-the-box solutions has to be completely rebuilt to exploit the new capabilities of the cloud. The majority of existing hardware-based playout solutions cannot be virtualised or used with standard x86 architecture CPUs as the majority of these solutions were GPU-based, making virtualisation challenging, and costly.

Important Trends Affecting Playout

Beyond technological challenges, the broader industry sees external shifts in perceptions of security and quality.

Content Security. Within the broadcast domain, digital media departments of broadcasters have already deployed video content to the cloud, eschewing some of the perceived security concerns linear broadcasters have been so focused on historically when evaluating cloud. In other areas post production SaaS solutions provide global platforms for collaborative content creation and media asset management. Leading public cloud providers such as Amazon Web Services invest heavily (compliant with all applicable MPAA infrastructure controls) and are committed to the highest security standards, which often cannot be achieved cost effectively even in private infrastructures.

Quality. On the other end of the spectrum, the consumer while still demanding high quality broadcast for premium content on a big screen, is getting accustomed to lower quality video-on-demand on the smaller screens (e.g. smartphones and tablets). Proliferation of YouTube and other often user generated but viral content has shifted focus from quality to speed of delivery.

With existing barriers to entry, one of the major ones being deployment of playout, how can content owners capture audiences not only through VOD but linear, often lucrative, distribution as well?

Cloud computing definition

The US National Institute of Standards and Technology defines cloud computing as a model for enabling “on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction”. Cloud computing resources encompass storage as well as networks, servers, applications, and services.

Cloud Infrastructure

Cloud infrastructure is comprised of both a physical layer and a software layer. The physical layer consists of hardware resources necessary to support the cloud services being provided, which typically include server, storage and network equipment. The software layer is deployed across the physical layer, which manifests the essential cloud characteristics by pooling and managing physical resources.

There are three main deployment models for the cloud, each with different pros and cons:

1. **Public cloud.** The cloud infrastructure exists on the premises of the cloud provider and computing resources are available for hire by the general public and corporates.
    a. Usually offers a wide range of related services, like petabyte-grade storage services or separate database services.
    b. Global technology leaders Amazon, Microsoft, Google, alongside smaller players, provide public cloud access, typically charging on pay-per-use basis.
2. **Private cloud.** The cloud infrastructure is provisioned for exclusive use by a single organisation. It may be owned, managed and operated by the owner, a third party or a combination, and can be located on and/or off premises.
    a. Servers and other hardware may be owned by the broadcaster and placed in its own data centre, or in a rented space of a third-party data centre.
    b. Building and owning a private cloud usually requires a substantial upfront financial commitment as well as on-going maintenance resources.
    c. Under certain conditions for broadcasters with very predictable and stable loads on computing resources, a private cloud may make economic sense.

3. **Hybrid cloud.** The cloud infrastructure combines two or more private and public cloud infrastructures, which are bound together by cloud technology that enables data and application portability. As an example, a broadcaster may archive all of his content in a private cloud block, but run the actual virtual playout servers from Microsoft Azure's public cloud.

Cloud Features and Benefits

To fully understand key cloud features and benefits, broadcasters should focus on:

1. **Accessibility.** To access and use cloud resources, no specialised, playout-specific, hardware is required. Computing resources can be remotely accessed through the Internet or a private network.
2. **Pooling computing resources.** Users benefit from the cloud provider's ability to pool computing resources, dynamically sharing physical and virtual resources. This pooling allows for cost reductions through more efficient load utilisation. Standards have evolved to ensure that pooling does not come at the expense of security. In fact, many mission-critical enterprise systems now run in the cloud.
3. **Elasticity.** A key feature of the cloud is its ability to scale rapidly in response to demand.

   a. In environments where there is substantial variability in demand for computing power (e.g. bursts when heavy encoding of video materials is required) elasticity is the key to reducing the demand on internal IT, which is impossible for traditional solutions to cope with.
   b. Also, if the launch of a new channel is unsuccessful, it can be shut down without incurring cost penalties, which is another clear benefit of elasticity.
4. **Redundancy.** Computing resources managed as part of cloud technology are combined in such a manner that ensures automatic back up of all resources at any point in time. By default, users can assume that media and other assets are automatically backed up and do not require additional solutions and resources.

Cloud Misconceptions

Virtualisation vs "True Cloud"

One of the misconceptions in the broadcast industry has been in defining the cloud as nothing more than a remotely accessed virtualised server in a data centre.

Virtualisation and the cloud are separate technologies, albeit with similar goals – to provide a way to run multiple software processes on the same hardware, thus increasing the mean-average hardware resources usable at the same time, and making sure that the available hardware is used to its full potential.

Virtualisation does not deliver elasticity and pooling comparable to cloud technology, and therefore demands substantial commitment of IT and engineering resources to manage virtualised applications.

Virtualising applications often originally developed for proprietary systems including GPUs is a challenging task and often requires specialised hardware/servers, which are not plain vanilla cloud resources. This approach does not allow users to extract maximum benefits from ever-falling cloud resource costs.

Remote Access vs. "True Cloud"

Some hardware manufacturers have been marketing remotely accessed traditional playout hardware servers as "cloud playout solutions", in some cases with additional layers of web-based consoles. The architecture suggests that you place the manufacturer's hardware in a remote data centre and operate playout remotely via the Internet.

In reality, such an approach has little to do with cloud philosophy. It has no elasticity because computing resources of such specialised playout hardware cannot be dynamically managed. Ultimately, users have to bear the high cost of acquiring and operating traditional hardware.

Service Models

It is important to understand key differences between service models of cloud based solutions:

1. **SaaS – Software-as-a-Service,** is an upper-tier level of service.
    - Client uses a software application supplied by a provider on a cloud infrastructure. For example, a cloud playout system like Veset Nimbus Playout, or a cloud-based content production system like Adobe Creative Cloud.
    - Software application is accessible from various client devices through a thin client interface, such as web-browser.
    - User does not manage or control the underlying cloud infrastructure, including the network, servers, operating systems or storage.
    - SaaS management places very minimal demands on internal IT resources.
2. **PaaS – Platform-as-a-Service.**
    - Users can deploy their own application (in-house or acquired) onto the provided cloud infrastructure.
    - The user does not manage or control the underlying cloud infrastructure, but does have control over the deployed application.
    - An example would be the ability to host a user's web application through the Amazon Elastic Beanstalk service. The user manages the web application while the service provider manages the actual web servers, including scaling, redundancy, database engines, etc.
    - For playout it would mean that the user (e.g. playout service provider) will develop his own, or acquire a license for, third party software and maintain/develop them internally for playout operations.

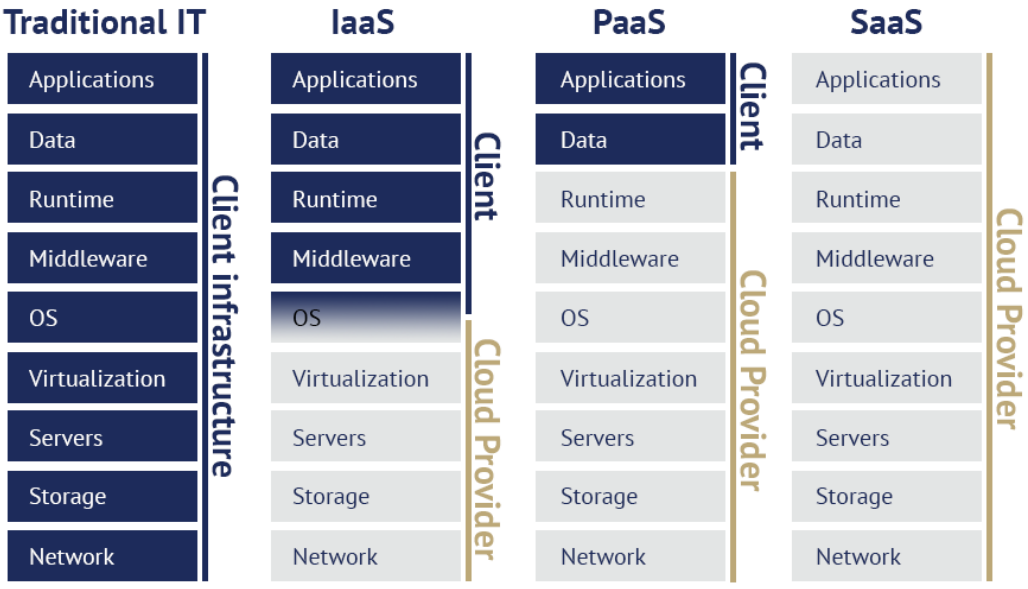

Figure 1

3. **IaaS – Infrastructure-as-a-Service.**
    - The user rents the infrastructure (processing, storage, network and other fundamental computing resources) to run his virtualised software and services, for example, virtual machines, database instances, etc.
    - The user does not manage or control the underlying cloud infrastructure, but does have control over operating systems, storage and deployed applications, and possibly some control over networking components (e.g. host firewalls).
    - While offering high levels of control, such a set up requires that substantial internal IT resources be committed to maintaining and running playout.

Cloud Architecture for Playout

So what is true cloud playout infrastructure and why is it important for broadcasters to understand and incorporate?

Service Model - SaaS

From broadcasters' perspectives, the most cost efficient and future-proof deployment and operation of cloud infrastructure is SaaS, for the following reasons:

- The broadcaster does not have to take care of the underlying cloud infrastructure – leading to reduced resources committed to IT maintenance;
- The playout service is rented on a monthly subscription basis with the freedom to increase, decrease or cancel playout services as required;
- Broadcasters can launch additional playout channels instantaneously (once MAM and transport is set up);
- The broadcaster pays for actual use of cloud resources during the period and no more.

What to Consider When Looking For Cloud Playout

Clearly, to use all the advantages of the available cloud infrastructure, playout should be developed as an independent software application that is not tied to proprietary hardware of a particular manufacturer or GPU/video card. Cloud resources run on various common operating platforms, Linux and Windows being standard.

Any substantial tie to specific hardware or software significantly reduces broadcasters' flexibility in the longer run and also may lead to higher costs. Furthermore, prices for public cloud resources have been falling consistently as leading operators reach economies of scale and pass those reductions through to their users.

Software can be more closely integrated into specific capabilities of various cloud services, such as Amazon Web Services or Microsoft Azure, two platforms that focus

on proving broader service offerings for broadcasters. However, it may nevertheless be wise to use multiple suppliers of cloud resources to improve redundancy and remove risks associated with being tied to a single provider of cloud infrastructure.

Selecting the right cloud playout architecture

There are multiple ways that a cloud playout system can be architected and deployed. The three main architectures are:

1. Edge,
2. Cloud, and
3. Hybrid

The main difference between these types is the location of the playout servers, and, in the hybrid cloud case, the location of asset storage as well.

Edge playout

Edge playout servers can be located in a broadcaster's private cloud (in-house or rented data centre), or more often at head-end (e.g. satellite, cable). Everything else (MAM, web access portal, ingest and transcoding) is run from the public cloud infrastructure. The Edge playout caches all of the assets, scheduling and titling information from the cloud service and outputs IP UDP (or ASI/SDI, with added hardware options) locally.

Edge architecture is currently the most popular cloud playout option mainly due to the fact that existing head-ends still cannot input IP signals and rely on ASI or SDI. Hence installing a physical server, usually a low specification one, is a workable, near-term solution.

Also Edge can offer savings on transport between playout and headend. However, the main disadvantage of the edge scenario is the hardware requirement, which brings the following challenges:

- Physical deployment – the broadcaster has to make often lengthy arrangements with the head-end owner to install such pieces of equipment;
- Edge servers require monitoring and maintenance both on site and remotely;
- Edge architecture increases complexity and time to air while simultaneously reducing the flexibility of the solution.

Figure 2

Cloud playout

All of the components of cloud playout – including the playout servers – are located in one or more public or private clouds. The IP output signal coming out from the playout encoder is transported to the head-end location either via the Internet by way of special IP transport protocol, or over a leased metro line. Such "pure" cloud playout can be deployed as SaaS and, if properly architected, should be cloud agnostic. This allows it to integrate with the software layer of any well-established cloud infrastructure.

The main advantages of having everything run out of the cloud:

- Fully centralised infrastructure, which allows full control over playout agents;
- The ability to launch channels rapidly;
- Far easier maintenance, no need to look after any physical infrastructure anywhere in the system;
- Use of fully redundant cloud servers (as opposed to physical for the edge);
- The ability to operate large number of channels / feeds globally (e.g. broadcasting channel to tens of head ends using any hardware infrastructure, such as edge, becomes a substantial burden).

The only meaningful downside of 100% cloud based playout is the need to transport feed from cloud to the head end, however, using modern IP transport, that can be mitigated.

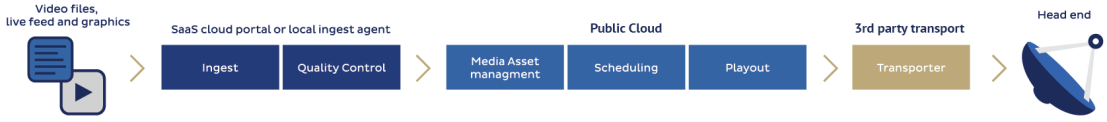

Figure 3

Hybrid playout

Hybrid playout suggests that the cloud playout architecture combines public and private infrastructures. For example the broadcaster operates playout servers from public cloud, but its central MAM is in a remote private cloud.

Figure 4

Transport

In Edge playout scenarios the encoded signal is output directly to the head-end's network and therefore directly to the multiplexer or further transcoders, as needed. If Cloud or Hybrid playout is used, the encoded signal is output on the virtual local network interface of the playout server. The signal then needs to be transported to the head-end. Because we are talking about linear playout, the signal must never be interrupted. There must be no frame drops. Traditionally, signals are converted to ASI-over-IP and transported using satellite links and metro lines. Those methods work well, but require additional physical convertors that cannot be virtualised. A further downside is that they are also rather expensive and provide almost no flexibility as satellite and leased line contracts are usually bandwidth-fixed and are signed for minimum of 12 months.

The Internet can be used to transport TV signals from playout to head-end, though special considerations must be taken into account to compensate for frequent jitter and latency in middle and last-mile connections. MPEG TS over UDP – the de facto protocol used to transport TV signals in private networks – does not work well in an Internet environment and will show considerable packet loss.

There are several solutions in the market that allow the use of the Internet for TV signal transport. They are built on either UDP or TCP; typically support error correction and a configurable signal delay with buffering; and can include additional useful functionality like one-to-many streaming, automatic failover switching, monitoring and transcoding, and recording. Some solutions require proprietary hardware equipment at both sending and receiving ends, thus failing both the flexibility and cloud philosophy requirements. Amongst pure software-based solutions, the ZiXi transport solution is mature, OS agnostic and works well in cloud scenarios.

The flexibility and cost efficiency of pure cloud playout, combined with a software-based, over-the-Internet transport solution, leads to new possibilities:

1. **Playout in full 1+1** mode – Each channel can be played simultaneously by two synchronised playout servers, each residing in separate public clouds. Signals from each playout server are delivered to the head-end. Failover switching in case of primary signal problems can be done automatically by either the transport solution (which ZiXi provides) or by the head-end IRD infrastructure.

2. **Custom feeds** for SD, HD, OTT, web – Cloud playout can be configured to output and encode custom versions of a signal for various distribution needs without any quality loss due to multiple transcoding. That is, a high bitrate HD signal with all available audio and subtitle tracks can be sent to satellite and cable head-ends; a lower bitrate HD signal to OTT distribution; and, finally, an SD signal with just one audio track to CDN for web streaming.

3. **Cloud-based monitoring and multi-view** – Monitoring is typically done by the head-end MCR (Master Control Room) team and requires special hardware-based solutions. A cloud-based playout and transport architecture can be further enhanced by a cost-effective cloud-based monitoring and multi-view solution, accessible in a secure way from anywhere with an Internet connection with, importantly, no additional investment in hardware and software licenses.

Coexistence with existing infrastructure and solutions

In conclusion, it is worth noting that cloud playout can easily coexist with traditional broadcast technology. Given its IT nature, cloud playout can be integrated with many aspects of existing hardware-based infrastructure broadcasters have already deployed, for example existing MAM, scheduling and other systems built into workflow. This allows a gradual replacement of legacy equipment and software which in turn enables new workflows to be implemented without major disruptions in operations.