PixelPower
A Rohde & Schwarz Company

Cloud Playout

Technology Evolution

Broadcast technology has now evolved to the point where it is possible to deliver linear and non-linear content from the public cloud. This white paper covers the following topics which merit consideration when evaluating cloud solutions for content delivery.

- How has technology evolved to facilitate content delivery from the cloud?
- What do we mean by a cloud deployment?
- What are the practicalities of moving to a cloud deployment?
- When and how and what should you move to the cloud?
- What are the advantages and disadvantages of a cloud deployment?

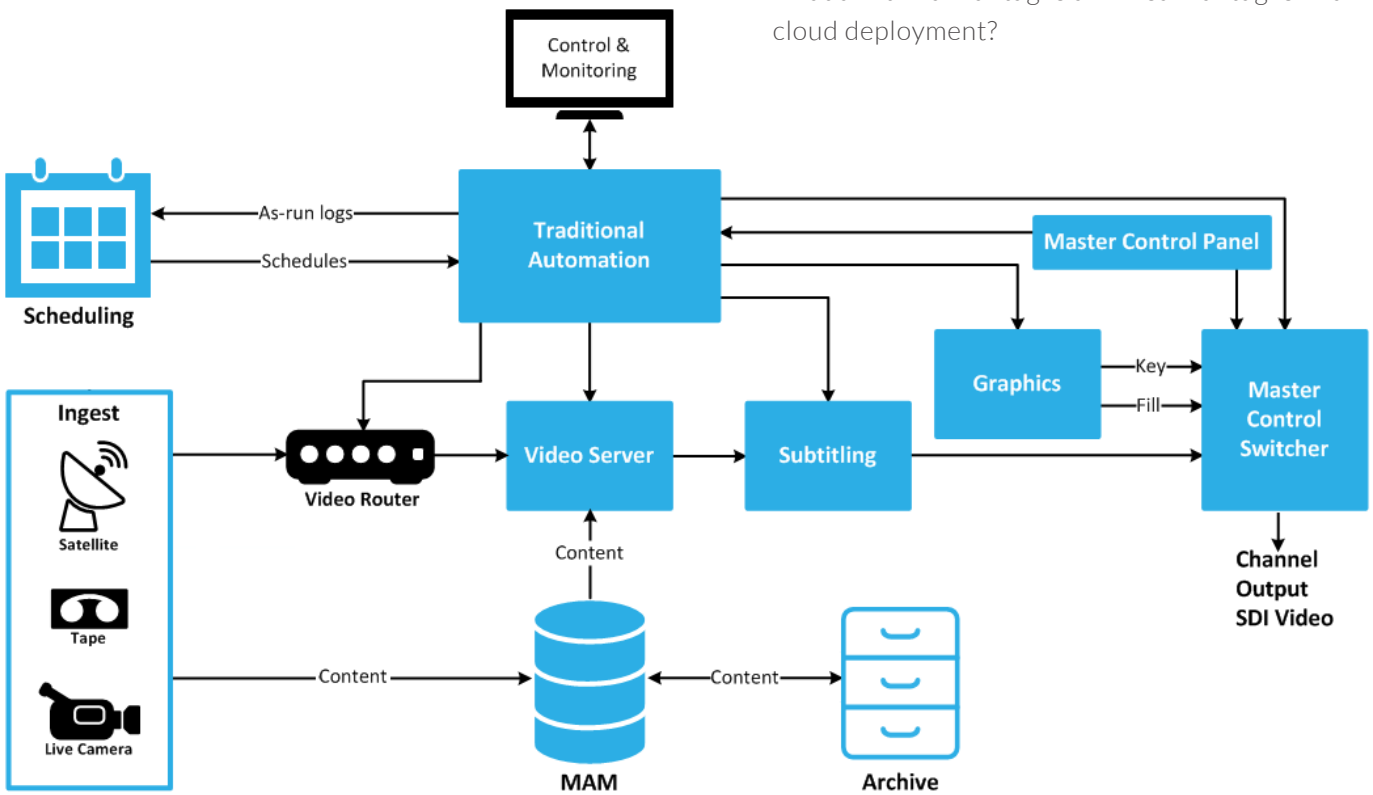

Traditional Playout Chain

Until quite recently a linear playout chain was built from separate pieces of dedicated hardware. General purpose computers were not powerful enough to execute the millions of operations required to process uncompressed video and audio streams in real time.

A broadcaster could choose the different components from a wide range of manufacturers – sometimes called a "best of breed" approach - and would expect the chosen automation system to control them and make them work together in a seamless manner to deliver their required output. Video and audio were connected together with SDI cable and control used serial or ethernet cabling.

A playout chain built in this way is expensive to buy, build and maintain – each supplier would have their own service agreement, and if there was a problem it could take some effort to identify the root cause and seek help from the correct supplier.

The configuration, being defined by physical hardware boxes and dedicated cabling interconnections, is neither flexible nor quick to change. If new functionality is required, a component would have to be replaced or a new component added, assuming the rack space and necessary power was available. A change of video format, for example moving from Standard Definition to High Definition, might require replacement of most of the playout chain.

Channel in a Box (or a few boxes)

As general purpose computing power increased and became more affordable it started to become feasible to implement certain functions of a traditional playout chain in software. This change enabled the introduction of "Channel in a Box" playout devices, which were able to collapse several functions into a single device.

However, some specialised functionality such as sophisticated graphics, DVE effects and complex subtitling requirements would still rely on dedicated hardware boxes. There are probably several reasons for this:

- The extra computation required for these functions could go beyond what the underlying general purpose computing platform could provide, leading to dropped frames and disturbance of the on-air signal.
- Good quality graphics and subtitling is built upon decades of experience and knowledge – by individual companies whose sole focus has been on evolving these best of breed solutions – so it is very difficult for a "channel in a box" vendor to develop this knowledge and experience when their history and expertise is more likely to be in video servers or automation.
- Graphics and subtitling rely on complex workflow processes to originate the content and these are likely to be deeply embedded and difficult to change. Keeping existing graphics and subtitling equipment in conjunction with an integrated playout device avoids having to change these operational workflows.

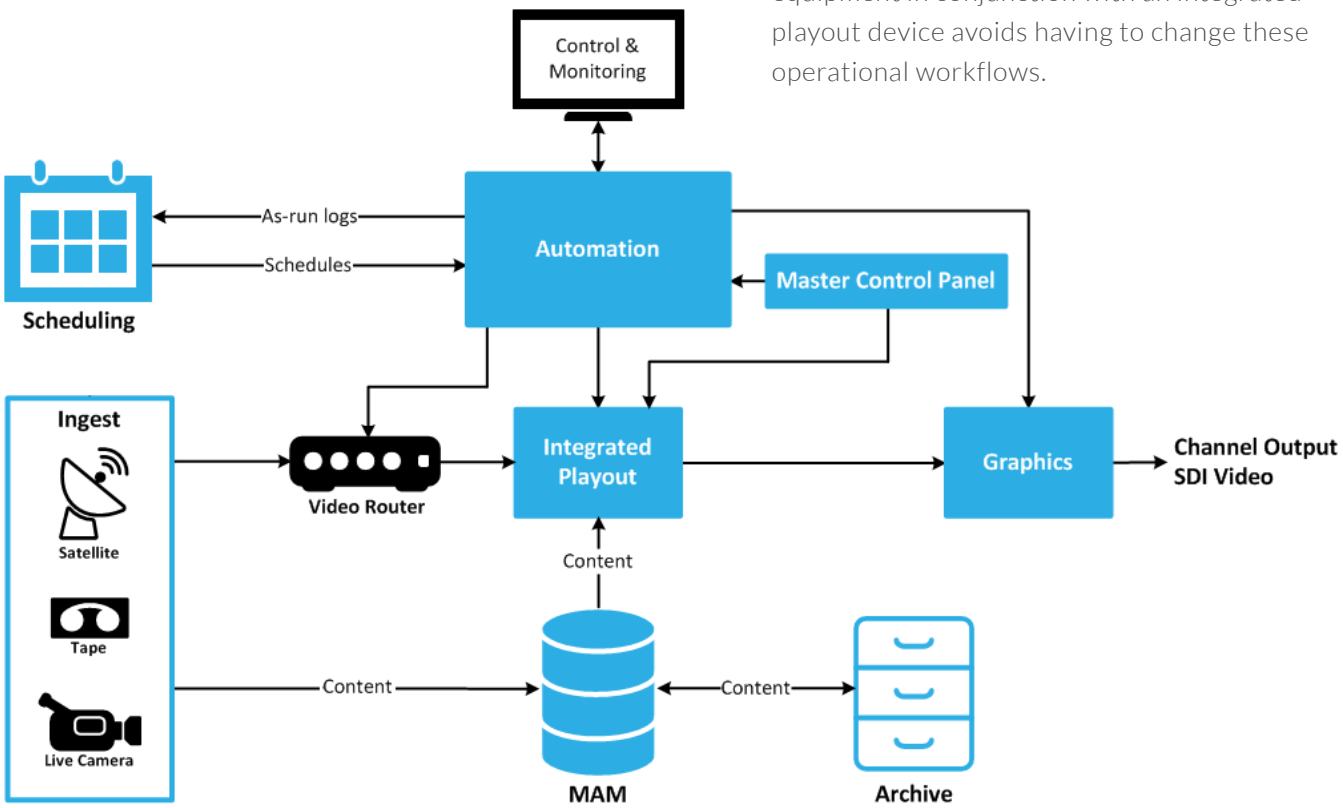

Software Defined Playout

General purpose computing power continues to increase and IP standards for video and audio streams are evolving fast.

New integrated playout vendors have emerged with deep experience in areas like graphics and subtitling. The hardware devices are replaced completely by software modules; SDI transports are replaced by IP streams. The dedicated hardware is gone and the complete linear playout chain is pure software. This means that it can be virtualized, and deployed in private data centre or even public cloud.

Further appetite for flexibility and cost saving amongst broadcasters and content owners starts to drive the adoption of software defined integrated playout solutions.

Software defined has many advantages over a discrete hardware approach, for example:

- It is more cost effective than the equivalent discrete hardware, both in terms of initial purchase price and in ongoing service and support.
- It offers the potential for opex (Pay As You Go) payment models which could be particularly suited to pop-up channels and Disaster Recovery scenarios.
- It is more flexible as broadcast standards evolve – a new IP codec or broadcast standard (e.g. UHD, HDR) can be implemented with a software change rather than replacement of multiple devices.
- It is more flexible if a new function is required; for example, audio watermarking can be added via a software update.
- It can support flexible deployment models – for example, a vendor might allow a software licence to be moved between locations or from bare metal to virtualized to cloud without any cost implications. This is not possible with discrete hardware.
- All the various software components operate seamlessly together, if something is not working as expected there is only one vendor to call (unless the hardware infrastructure or hypervisor is at fault).
- It can still offer some of the benefits of "best of breed" approach by integrating third party plug-ins for specialist functions such as graphics or audio loudness processing.

There is perhaps one disadvantage with software defined; it is less deterministic in its performance than dedicated hardware performing the same functions. This can be mitigated by careful configuration and use of tools provided by the vendor for evaluating and monitoring performance. It is important to match the underlying off the shelf computing hardware performance with the expected requirements and be prepared to upgrade the (relatively low cost) COTS hardware if significant new functionality has to be added in the future.

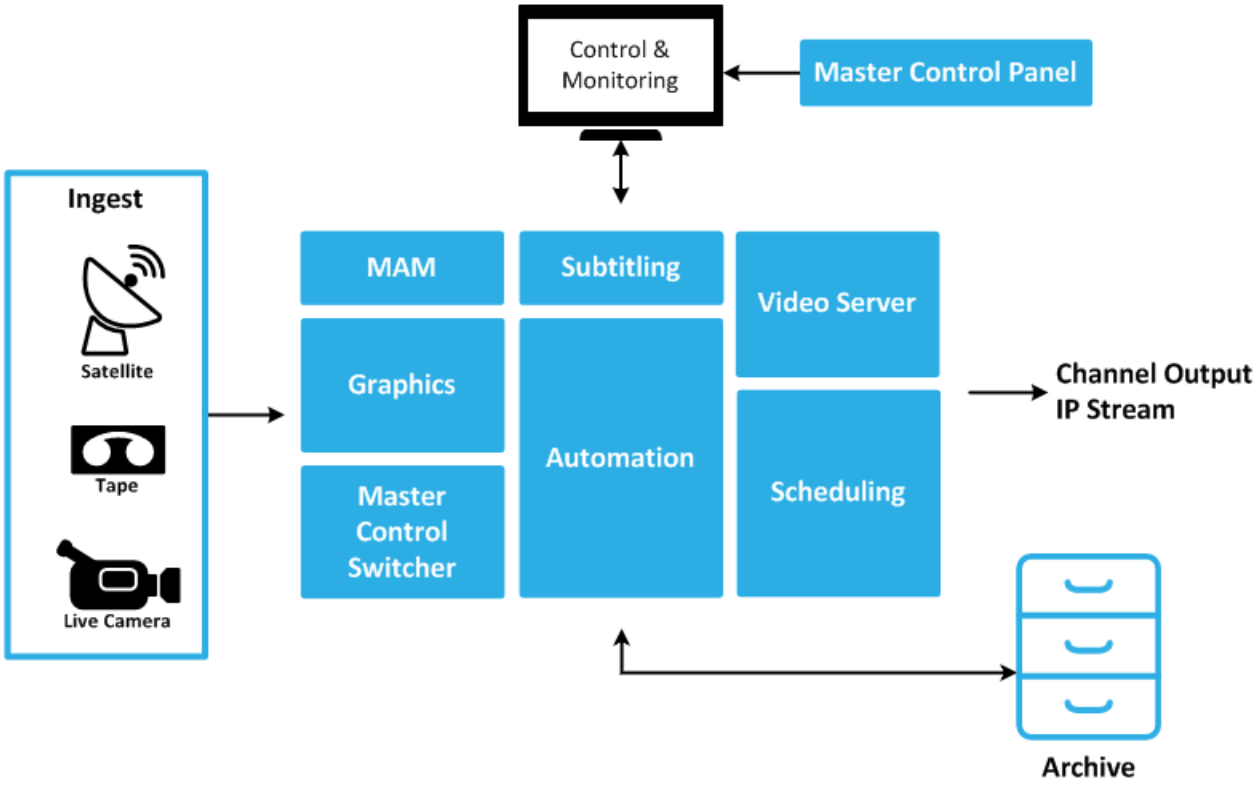

What do we mean by cloud deployment?

"Cloud" means different things to different people. For some, it is a private data centre; for others, a public cloud provider such as AWS, Azure or Google. Or simply "someone else's computer".

Cloud implies a degree of flexibility in terms of how the requirements are matched to the infrastructure and almost certainly implies virtualized solutions and remote operation. Obviously the chosen vendor must guarantee their solution is capable of operating in a virtualized environment and can be controlled and monitored remotely.

A private data centre (whether owned by the broadcaster or a third party service provider) invariably offers more flexibility in infrastructure than a public cloud provider can offer. For example, it may be possible to install dedicated hardware equipment (for example, encoders and multiplexors, uplink equipment) or even some SDI infrastructure if needed. In many ways a private data centre is not much different from a broadcasters own equipment rooms – apart from the prevalence of virtualization and the ability to monitor and operate the solution remotely.

For a public cloud implementation the situation is rather different; there is no possibility of changing the infrastructure so the solution must be able to work within the constraints of the specific cloud providers' offering. Those constraints are likely to include:

- Network does not support multicast or broadcast packets (may preclude use of PTP timing)
- Not feasible to use uncompressed video processing and routing
- No guaranteed uptime or Service Level Agreement

From this point on we will be focussing solely on public cloud, rather than private data centre deployment.

Why might you want to use public cloud (or not)?

If you've got this far, you are likely to be seriously considering deploying some or all of your content delivery to the public cloud.

There are many reasons why you might want to do this, for example:

- It can be purchased on a usage basis; there is no need to invest up-front in hardware infrastructure.
- Content is delivered to the consumer via CDN / public internet, so originating it in the cloud makes sense.
- The capability is not required all the time, so it doesn't make sense to pay for it all the time, for example Disaster Recovery or Pop-up channels around festivals or sporting events.
- It is scalable up and down – for example, if new capabilities are required additional cloud infrastructure can be added, and then shut down when no longer required.
- It is flexible in capabilities – most public cloud vendors offer a range of different virtual machine types with varying degrees of computing power. If more features need to be added to a software defined playout then it can be launched on a more powerful virtual instance, without having to replace physical hardware.
- It is flexible in geography – most public cloud vendors offer a wide range of locations, which may be useful for legal and commercial reasons (where data is stored, where taxes are paid, where a more favourable regulatory regime is in place) and practical ones – regionalisation to originate the content more locally whilst still addressing a global marketplace.
- Changes like adding functionality to a channel configuration, starting and stopping new services can all be made very quickly, often in less than a day. This is not possible with fixed infrastructure and traditional playout.
- It is a potentially seismic change to operational practice which may present an opportunity to streamline other parts of the operation even if they are not moving to the cloud.

On the other hand, there are some disadvantages. Probably the biggest one is cost – certainly for 24/7 channels a cloud deployment will cost several times more over a 3 year period than hosting the equivalent infrastructure on your own premises. It is also not that easy to calculate the cost of using public cloud – some costs are based on data throughput, some costs are purely based upon time usage and others are based upon number of accesses to a particular function. There are other reasons why you might decide not to use public cloud, for example:

- You rely on specific functionality in your playout chain which is not available for use in the public cloud due to architectural or licencing constraints.
- Your preferred vendor does not support cloud deployment of their solutions.
- You don't want to be "locked in" to a single vendor.
- Your team does not have the necessary skills and knowledge to deploy and operate a cloud based solution.
- Legal restrictions prevent you from deploying outside your country and your preferred public cloud provider does not have a data centre in your country.
- You have concerns over the security and control of content and customer information when deployed in the public cloud.
- Public cloud infrastructure is optimised for traditional IT tasks, not real-time processing of video, audio and metadata.

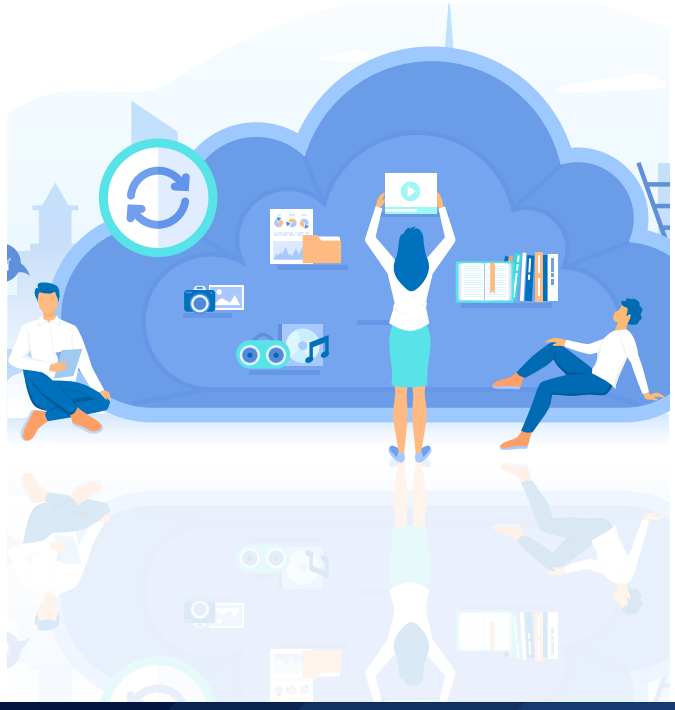

What do you need to consider when deploying in the public cloud?

Having considered the pros and cons of public cloud deployment, we now consider the practicalities of implementing a cloud based playout. These are broadly divided into Commercial and Technical considerations.

Commercial

Choice of Cloud Provider

Choosing a cloud provider is also an important part of the process of moving to the public cloud. Although there are many cloud service providers, some may be eliminated for the following reasons:

- They do not offer a data centre in the regions where you need one.
- They may not be supported or recommended by your chosen playout solution provider.
- They may not provide the capabilities that you need for your implementation (for example, no direct connection in your city or no GPU equipped virtual machines in your region).

Remember that the larger cloud providers are targeting the general IT/e-commerce marketplace, replacing traditional business back-office systems rather than working in the highly specialised world of broadcast. In the past they have had very little domain knowledge so you needed to rely on your chosen solution vendor to guide you; that situation has changed to some extent, with many of the major players having their own specialist media teams to advise.

One area that will need careful attention when comparing different offerings is the costings – as already mentioned, cloud providers have complex pricing models with a myriad of services and add-ons available. Detailed financial modelling may be necessary to find the most cost-effective platform on which to deploy. Part of the cost equation is also about changing cloud provider, should you need to do that for commercial or political reasons. There might be substantial costs involved in moving your archive between providers. One possible solution to this is to adopt a cloud provider agnostic approach – many third party companies offer orchestration layers which offer a single "virtual" cloud, allowing relatively easy switching between underlying cloud service providers. This gives more flexibility in changing suppliers as the cloud infrastructure marketplace becomes more competitive.

Commercial Models

There are different ways of using public cloud infrastructure for your playout; broadly speaking, you can treat it as IAAS (Infrastructure as a service), renting the infrastructure from the cloud provider and being responsible for software installation and operation yourself; or you can purchase a fully managed service through a third party vendor who themselves purchases the infrastructure from the cloud provider.

Many traditional playout service providers and also newcomers to the market offer managed cloud based services and they can provide a simpler, albeit less cost-effective route into cloud based playout. But for optimum cost reduction and greatest flexibility and control, running your own cloud based infrastructure is preferred, provided you have the skills (perhaps in partnership with your chosen vendor) to build and manage it.

Security and Skills

Deploying to and operating cloud infrastructure requires some specific new skills which may not be available in the organisation – it may be possible to upskill existing staff, or use external help (a consultant, system integrator or the solution vendor) to assist with the migration process.

Another concern is that of security – you will be putting valuable content and probably customer information covered by GDPR in the hands of an external company, rather than keeping it in-house on your own servers. But the reality is that the leading cloud providers have very robust security and encourage their customers to use it. The security team at even a small cloud provider have a skill set and experience far greater than that of even a major national broadcaster and there should be comfort in that. Broadcasters might typically spend 10-20% of their operational costs on security for in-house systems and this saving should be taking into account when comparing with cloud based providers.

Putting playout in the cloud also mandates clear separation from office based networks, reducing the risk of security breaches. Also, most security issues are caused by human error, so the increasing automation that a cloud deployment enforces will reduce the chance of that.

Optimise Workflows

Moving to a cloud playout represents an opportunity, or could even make it a necessity, to optimise and improve workflows and processes within the organisation. There is little chance of customising the infrastructure in the cloud so this can be used to help eliminate unusual workflows (no more "because we've always done it like that") and accept more standardised practices as supported by the chosen vendor's solution. Standardisation and accepting offered solutions without requiring expensive customisation will also help to keep costs down

Technical

Content

If your content archive is not already stored in the cloud then it will need to be uploaded; many cloud providers offer one-off data migration services using the cheapest available high bandwidth connection – using a standard carrier to transport high capacity disk drives between locations.

Once the archive is available within the cloud infrastructure you will also need enough upload connectivity to keep it updated with new material, and potentially provide live contribution feeds such as from studios or outside events. Many cloud providers offer "Direct Connect" or equivalent, which is a dedicated connection from a local POP (point of presence) to the cloud data centre. This can give higher guaranteed bandwidth connectivity but is no guarantee against packet loss. For live streams forward error correction will still be needed for contribution feeds or bringing down the channel output feed for further distribution (for example by local cable or satellite uplink). Forward error correction is also important for bringing monitoring feeds back to the point of control; there are many commercial solutions available, for example Zixi, RelayCaster from Rohde & Schwarz and the open-source SRT (Secure Reliable Transport) protocol promoted by the SRT Alliance. Some POP's (particularly those used by broadcasters and distribution operators) are gearing up for broadcast traffic and may offer specific provision for reliable delivery of real-time video feeds.

Monitoring and Control

Monitoring and control also needs careful consideration. The automation clients could be cloud hosted, with simple remote desktop access from any location to interact with them. This simplifies the security (by fully utilising the cloud providers' security models) and allows operation from any location. Alternatively, the clients could be installed locally, with appropriate connectivity through a firewall to the cloud playout infrastructure. On-premise clients potentially give a more fluid user experience, which might be important for live event switching.

Monitoring in a traditional master control playout would normally use multi-viewers showing the various incoming and outgoing feeds. There might also be signal integrity checking, either built into the multi-viewer or as a separate device. There is the same choice for monitoring as for the automation control clients; the monitoring could be hosted in the cloud and accessed and viewed via remote desktop technology, or it could be local with the required feeds brought down for monitoring.
If the distribution of the channels is not directly via CDN/internet (e.g. satellite uplink, cable, terrestrial transmitter) then the live streams need to be brought back to the uplink point anyway, which might be a suitable location for local monitoring and control. For internet distribution it probably makes more sense to keep the monitoring and control fully cloud based.

Location

One decision that requires both commercial and technical input is the choice of data centre location. There may be legal or regulatory constraints to where the cloud playout can be situated, but there are also technical constraints such as latency and the availability (or not) of high bandwidth connectivity. There may also be significant cost differences between different locations – it may be cheaper to use a data centre on the other side of the world if the latency can be managed. Don't forget to consider a hybrid location model as well – for example, host 24/7 channels in your own equipment room, but locate Disaster Recovery and Pop-up channel capacity in the public cloud.

Standards

Choosing the right IP standards for your cloud playout is also a challenge – the standardisation process often lags behind vendor technology offerings, leading to the risk of adopting something which later becomes non-standard or obsolete. If you are relying on a single vendor this might not be a problem in the short term, but it can reduce flexibility for the future. A lot of current standardisation effort is around ST-2110 uncompressed video but this is not generally a viable option for cloud based playout; cloud network infrastructure is unlikely to support the required bandwidth and packet throughput and the egress cost of uncompressed video would be prohibitive. A more likely mezzanine format of 30-50Mbit H264 RTP MPEG-TS is well established and a better fit for the cloud infrastructure and use cases.

Latency

Another important factor is latency. There are three main contributors to latency. When a signal is encoded or decoded (for streaming or display) a few frames delay (dependent on encoding type and codec parameters) will be introduced. When a signal is sent up or brought down from the cloud the necessary forward error correction will introduce delay; forward error correction latency is typically configured as ten times the round trip ping time of the location, so can vary from under a second for an in-country data centre to several seconds for one elsewhere in the world. The third source of latency is in the control and orchestration layer; when an operator presses a button, how long before that is actioned; this is also related to the round trip ping time to the data centre. When these latencies are added together there can be a considerable lag between an operator pressing a button (for example, to take out of a live event into a commercial break) and the action being seen on the monitoring stream. The monitoring and control should be designed to minimise the latency in order to meet the operational requirements of the channel; there will also be some degree of adjustment required for the operators when controlling a cloud based playout with its inherent latency.

Redundancy

A traditional playout would typically have a degree of redundancy, so that if any component fails then a channel will not go off-air. For tier 1 channels this would typically be a fully redundant 1+1 (or even 1+1+1) configuration; for less important channel groups an N+1 model might be more likely. Redundancy is also essential to allow a system to be upgraded whilst it is still on-air.

These traditional playout redundancy models can also be applied in the cloud if appropriate, but there are other things to consider:

- If the cloud playout is acting as a backup (+1) for a channel located elsewhere then it may not need to be inherently redundant itself.
- For implementing 1+1 redundancy in a cloud infrastructure consider locating each chain in a different data centre or at least a different availability zone to give further resilience. These options are more resilient than a 1+1 implementation in your own rack room or data centre because they are more physically separated.
- If the cloud playout is for disaster recovery purposes then 1+1 redundancy (possibly across two data centre locations or availability groups) is desirable. The increased costs of 1+1 are only incurred if the system is used.
- For upgrading a cloud based playout the new version could be installed on new virtual machines and fully tested in parallel with the on-air system before switching over and shutting down the previous version. The upgrade process takes advantage of the inherent benefits of cloud infrastructure and does not require traditional redundant infrastructure.

Integration with other Solutions

It is unlikely that one provider will give everything you need, and in any case it might be wise to choose other parts of the overall solution from alternative vendors. Integration might be required with a scheduling system, a MAM, a transcode farm or a CDN/VOD platform. The ability to draw upon significant cloud computing power at will might support AI add-ons, for example for automatically adding metadata to video and audio content in a media library, or selecting content based upon audience big data.

Whatever the integration needed, the chosen supplier must be able to support licencing and implementation on your chosen cloud platform, potentially for both continuous and occasional usage (pay as you go).

Availability

When you have your own equipment in your own rack room its availability is clear – if you have power, air conditioning and working equipment then channels will play out. The picture is not so clear with cloud based infrastructure.

The first availability question is when you want to start up some infrastructure, for example in a disaster recovery situation or for an event based pop-up channel. Does the cloud provider have sufficient capacity to give you what you need to run the channels? Availability might not be guaranteed unless you pay for it, perhaps by reserving critical parts of your infrastructure to guarantee availability. In a disaster recovery situation it is possible that other local broadcasters or businesses might use the same data centre, so if a natural disaster happens the data centre spare capacity is likely to be used up quickly.

Another question is your connectivity to the cloud infrastructure, whether through direct connect or public internet. How guaranteed is that connection through your local ISP? What SLA do they offer? You should have at least two independent connections (not going through the same trench or conduit) for operating your cloud playout. These two connections should ideally be from different locations, so your operations can be moved if one location is unavailable or the connectivity is lost.

Finally, there is the general reliability of the cloud providers' infrastructure. Do they offer an SLA? What uptime do they guarantee? The answers to these questions will inform the decision on what kind of redundancy you need for your cloud based infrastructure.

Case Studies

Up to now we have being talking about the general theory and practice of cloud playout. Now we will look at a couple of real world examples and outline the architecture in more detail.

Disaster Recovery for Tier 1 Middle Eastern Broadcaster

This system was designed to provide a low cost disaster recovery concept for a top Middle Eastern Broadcaster. Whilst they had redundancy in their on-site playout they had no backup if their building was disabled through power loss, fire or other physical disaster.

The solution was built on AWS outside their region and was designed to have minimal operational cost when not in use.

Content from their staging MAM and the upcoming schedules were continuously synced 7 days ahead to S3 cloud storage. Limiting this to a week kept the S3 storage costs low; this was the only operational cost when not in use.

When the DR system was needed, either in an emergency or for planned maintenance, the Pixel Power Gallium Automation and StreamMaster Integrated Playout virtual machines were spun up and could be playing out content in sync with the main facility within 15 minutes from launch. When running in parallel the live schedule from the on premise automation was continuously synced to the cloud playout, which would stay in line with any changes made on the main playout system. Once the main playout system was no longer available the cloud playout would continue independently, with a buffer of 7 days of material. A week was considered a long enough period to either recover the main playout or put more content and schedules in place on cloud storage.

The cloud based DR system could be operated from a laptop using the Gallium automation thick client, from any location in the world. The linear channel playout was brought down to the uplink site over public internet, with suitable forward error correction in place. The channel could also be served OTT via CDN.

The cloud infrastructure cost when the system was running was approximately 10 USD an hour (EC2 instances hourly usage charge plus data out costs, does not include Pixel Power software licence costs).

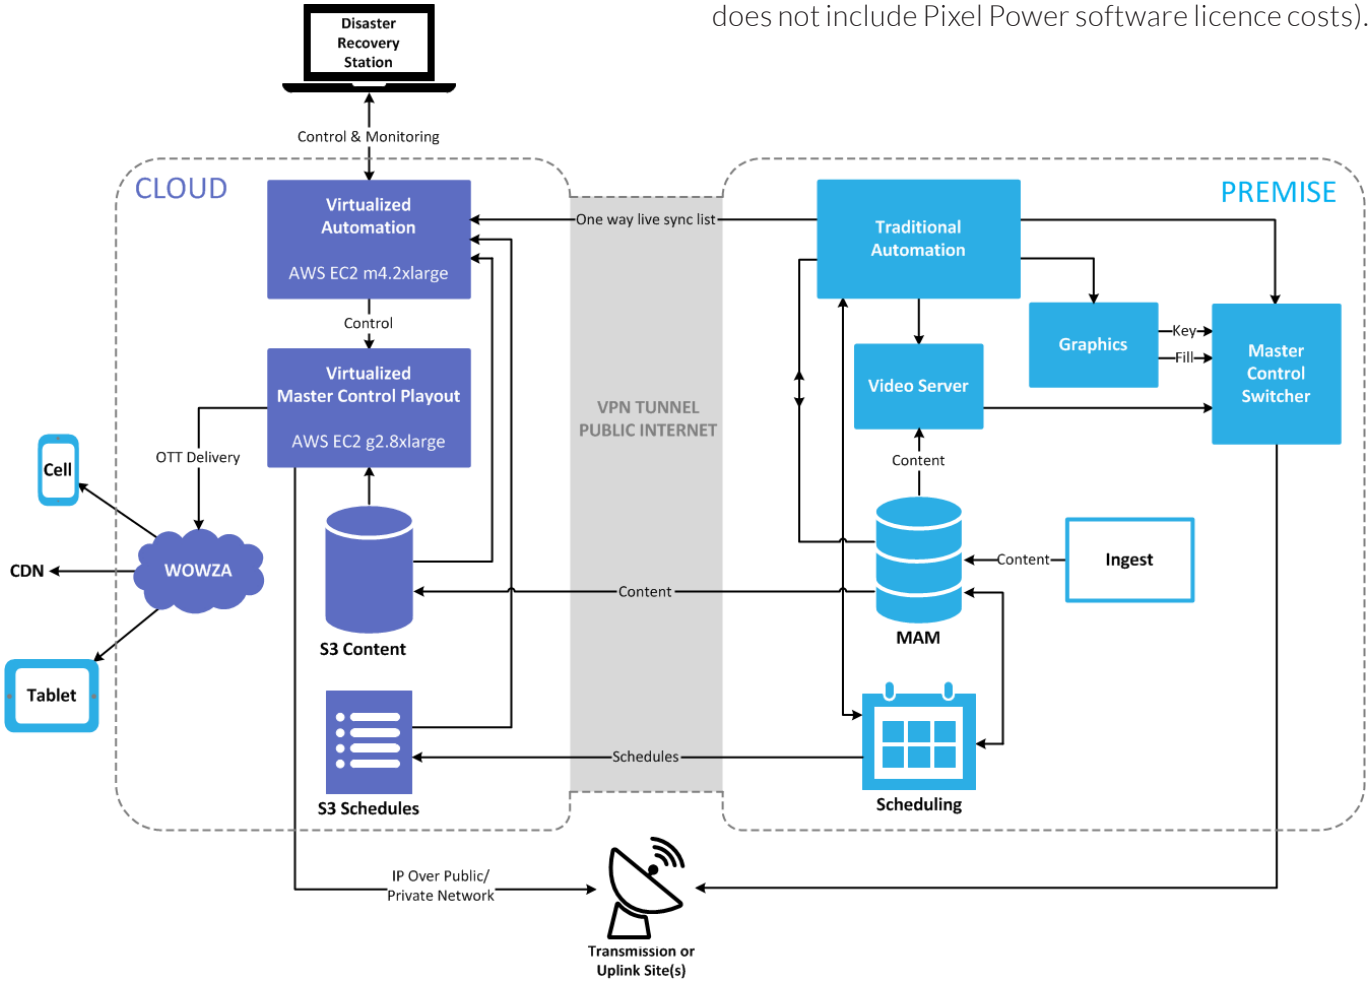

Thematic Channel Playout for Tier 1 European Broadcaster

This system was designed to show how a cloud based playout could support the future needs of a tier 1 European Broadcaster for its thematic channels. This was designed for primary playout, not as a Disaster Recovery system. The requirement included subtitling, Dolby audio processing and sophisticated graphics, including 3D DVE moves along with channel variants including SD simulcast and opt-out channels with alternative content and geo-blocking. A single Gallium automation virtual instance was configured, along with the StreamMaster instances to deliver the 9 required feeds.

The cloud infrastructure was built on AWS and used the broadcasters' direct connection pipe for control and monitoring. The MWEdge encoder was used for forward error connection on the mezzanine streams, which were downlinked as 30Mb H264 encoded MPEG TS and converted back to SDI at the customer premises for high quality monitoring on a broadcast monitor. The content was already available on S3 storage.

The software defined multi-viewer was run on AWS and its combined output brought back via the same MWEdge encoder/decoder infrastructure. For the Gallium client both an HTML5 remote desktop client (using low latency gaming technology) and a local thick client (using direct connect to the Gallium automation server) were evaluated. There was not much to choose between them from an operational point of view – some operators preferred the remote desktop client, others the more fluid experience of the local thick client.

The redundancy model for the test concept was N+1; for full deployment the infrastructure will be doubled up to give full 1+1 with a +1 guard on each side.

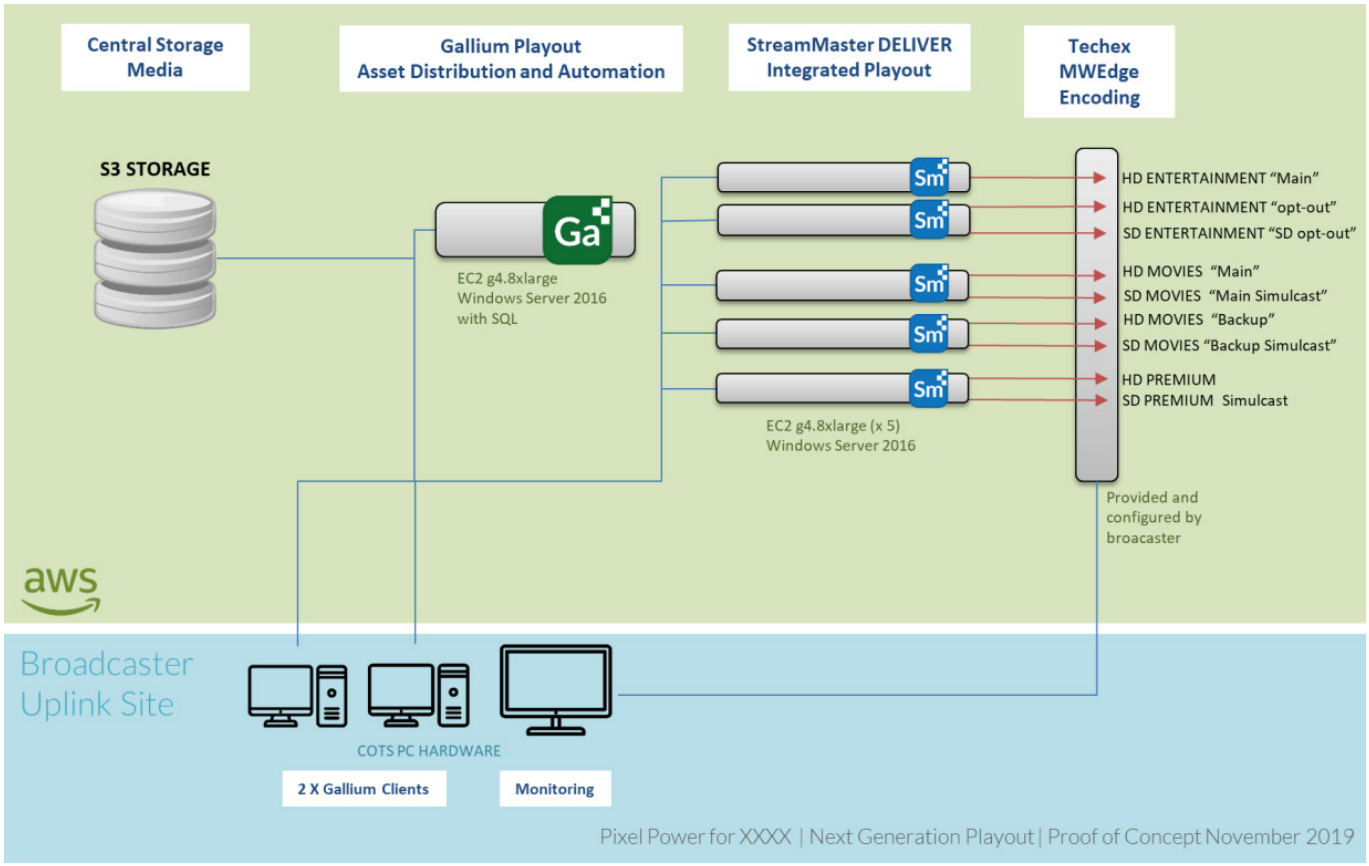

Conclusion

Moving to cloud playout is a challenging thought and there are indeed multiple issues to consider. It's a big decision and the buzz in our industry over several years tends to excite us in the direction of wondering whether it is worth it.

Is it possible? Yes, for sure, it's possible and there are now quite straightforward ways to make such a move happen. But as we know from experience with broadcasting over many decades, while our industry is filled with broadcasters, each and every one has a very different process and way of working. The considerations may be wide and varied but their path is now clear and can be systematically defined with a carefully thought through plan. Some years ago a similar decision was encountered when the industry moved from tape based workflows to file based workflows. All content, current and historical had to move to digital files. The flexibility this offered came to outweigh the drawbacks in time, effort and cost of digitising vast amounts of content.

In the move to cloud playout there is a very worthwhile stepping stone in the process that helps bridge what could, for some, be too big a leap. Moving first to virtualized software defined workflows enables broadcasters to take advantage of the more immediate benefits of remote connectivity, flexibility and future-proofing that software only applications can offer. Moving to the cloud could be considered a separate step, following the mastery of virtualized workflows - the workflows are relatively straightforward to flip into the cloud leaving the remaining challenge of taking that tape library to its final destination in the cloud along with current content.

We hope the information provided in this white paper will be helpful in understanding some of the things you need to consider in moving to a cloud playout solution; whether you decide to make use of cloud will ultimately depend very much on your own needs and constraints. Every use case is different, but having a good understanding of some of the pitfalls, challenges and benefits should help you make a more informed decision.

UK, Head Office
Pixel Power Ltd
College Business Park
Coldhams Lane
Cambridge
CB1 3HD
UK

Tel: +44 (0) 1223 721000

USA
Pixel Power Inc.
13457 Colfax Highway
Grass Valley
CA 95945
USA

Tel: (818) 276-4515

Middle East
Pixel Power Ltd
Dubai Studio City
Building 4,
Ground Floor, Office 2
PO Box 502459
Dubai
UAE

Tel: +971 (0) 4 4329282

Singapore
Rohde & Schwarz Regional Headquarters
Singapore Pte. Ltd.
(Pixel Power Ltd)
9 Changi Park Vista,
#03-01
Singapore 486041

Tel: +65 6307 0810

sales@pixelpower.com

Version 1.2 July 2020

PixelPower
A Rohde&Schwarz Company