SET Brazil

SINCE 1916

SMPTE Motion Imaging Journal

Covering emerging technologies in film, broadcast, and the connected media ecosystem

Media Cloud Fundamentals

By Tomasz Witkowski and Richard Welsh

Abstract

This paper aims to provide an overview to the reader of the current state of cloud use in the media industry, the adoption in specific areas, and the challenges facing current workflows. It further discusses specific areas that will provide major shifts in cloud use and have a significant impact on the media industry. The flow of use across all stages of production, live broadcast, mastering, versioning, distribution, and analytics is covered. This paper assumes a basic understanding and familiarity with the principles of industrial/enterprise cloud, software defined resources (SDx) and content delivery networks (CDN).

Keywords

Analytics, artificial intelligence, big data, cloud, content delivery network, distribution, live broadcast, mastering, media life cycle, packaged media, post production, pre-production, production, versioning, workflow

Introduction

This paper explores the basics of the current state of the art in media cloud workflows, where the challenges lie, and what the future holds for the media industry when it comes to cloud. We are skipping "cloud 101," so we are assuming that the reader is familiar with the following concepts:

- *Public Versus Private Cloud*: A commodity main cloud provider such as Amazon Web Services (AWS) or Microsoft Azure versus a privately owned and run cloud built on a framework such as OpenStack or Xen.
- *Infrastructure as a Service (IaaS)*: The underlying hardware that constitutes cloud, in principle, software-defined network, compute, and storage in one or more data centers. This fabric is invisible to the end user, presented through a service layer, the platform.
- *Platform as a Service (PaaS)*: The actual interface to the cloud infrastructure exposed to customers. In practice, this layer is used by enterprise customers and cloud service providers, and it is typically the domain of the developer. This layer has all the native services available through the cloud, such as Azure Blob storage or AWS Elastic Cloud Compute. These services are leveraged by service providers to build their software on, which is presented to the end user.
- *Software as a Service (SaaS)*: This is the layer most users and consumers of cloud services will see. This is the abstraction of the service from the hardware required to support, and it can be presented via user interfaces, application programming interface (API), or to back out a hardware-supported service such as Alexa or Apple TV. It should be noted that, in the SaaS layer, a software platform may exist which is distinct from the PaaS layer that is only concerned with the services native to that cloud, such as Netflix, Apple Music, or International Business Machines Corporation (IBM) Watson.
- *Cloud Native Versus Forklift*: Cloud native refers to SaaS that is built on the layers described above. Forklift (or lift and shift) refers to taking software that is designed to run on a local machine and moving it wholesale to a virtual machine (VM) in the cloud. The former is a fully abstracted service where scale is automatically defined by the workload presented to it. The latter is virtualization of existing hardware units where scale is manually defined by the number of VMs in use.

This paper explores the basics of the current state of the art in media cloud workflows, where the challenges lie, and what the future holds for the media industry when it comes to cloud.

Now we have got that out of the way, let us get into the good stuff!

Ideal Media Cloud

Recent tools and technology allow the entire media life cycle to tap into the cloud as a central resource. Cloud is more mature in some areas than others, and there are certainly challenges in some areas as will be discussed later. However, in principle, it is possible to do everything in the traditional media life cycle in the cloud today. The content flow for an ideal media cloud life cycle is shown

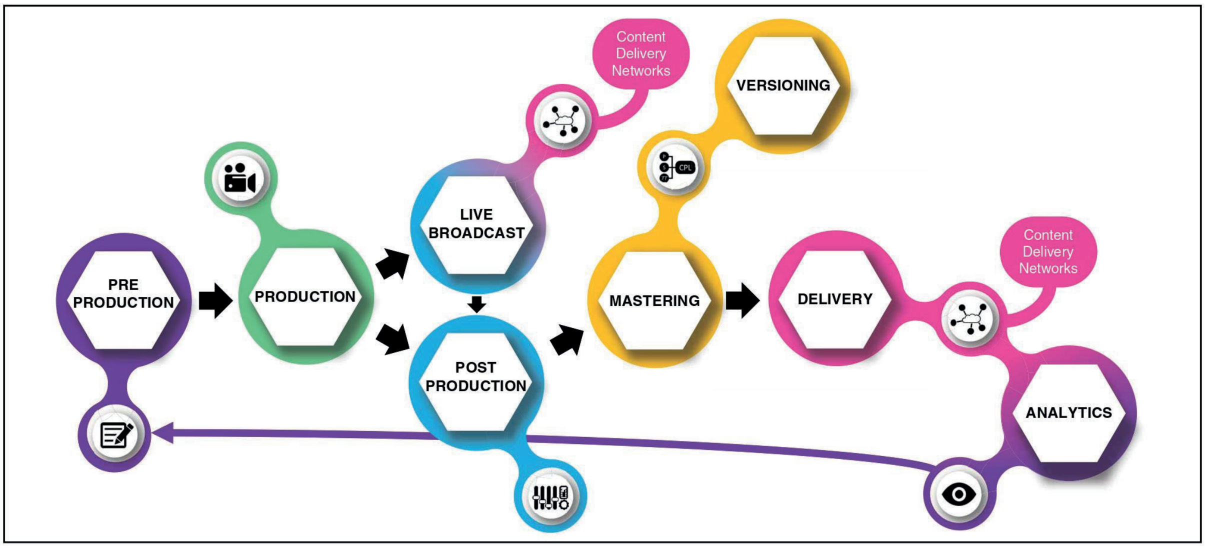

FIGURE 1. The idealized content life cycle in the cloud.

in **Fig. 1**, where all media data is centrally available and applications at every stage come to the data in situ. The data is only decentralized at the distribution stage when media is published to content delivery networks.

Pre-Production

In the development, scripting, design, previsualization, and other preparation for all types of production, the cloud is already well adopted. This is primarily in the area of collaboration on pre-production materials that are easily shared and kept as live working documents through a number of low cost or free everyday cloud services such as Google Docs, Dropbox, Microsoft OneDrive, and so on. Other collaborative tools hailing from the world of software development, such as Slack, have also become commonplace for realtime collaboration. However, this has extended more recently to direct content creation through previsualization, pre-production video such as location scouting, costumes, makeup, visual effects (VFX) design, and continuity, and now contextual metadata generation such as taxonomy and character information. Next-generation cloud-enabled tools in this space will include immersive (for virtual set walk through, next-generation location scouting, and practical effect design), artificial intelligence (AI) for scripting, plot development, and analytical feedback from previous related productions.

Production

Cloud services are widely used, again generally in the collaboration space. Dailies are probably the most prolific use case where there are multiple options and very well-established tools that cater for the varying needs of more traditional film production through to the rapid pace of reality TV and weekly episodic shows. They typically reach beyond sharing of video to include call sheets, continuity, security controls, and so on. Extending this ability to quickly review, remote live review via the cloud is now also an established space allowing producers to review the work on set and the directors to collaborate with second units all in realtime.

Primary workflows (camera RAW data and metadata) still largely employ on-set hardware and offline delivery (RAW data may well be sent to post via secure file transport but this is often a separate step in the process, for instance, overnight upload of most important takes on a shoot). In live, realtime movement of camera data to the cloud is more common but this will usually be a compressed contribution format. There are now bonded cellular solutions that allow direct camera-to-cloud and digital imaging technicians (DITs) to upload footage in remote locations without physical connectivity. This is still not commonplace but is increasing in popularity, particularly as costs drop. With further rollout of 5G networks, this area is likely to take a quantum leap and may

> **The origin of forklift**
>
> As a side note, forklift originated from large-scale IT infrastructure changes that literally required pallet loads of hardware and, thus, a forklift truck. Later, this came to mean any hardware or software upgrade to IT infrastructure that requires large amounts of effort on the part of the IT department. With the cloud, it seems to have been adopted, now taking on the implications of lifting the existing software from the ground into a VM in the cloud and forking the software development of existing applications to have a VM specific branch (hence fork + lift). Cloud natives, stretching the cloud–sky analogy, tend to refer to software running on local machines as being on the ground.

supplant the use of traditional high capacity landlines at facilities or physical disk transport solutions such as Amazon's Snowball.

Live Broadcast

Live is a rapidly growing area in the cloud. Most established streaming services now have live events, particularly sport, and in the case of traditional broadcasters supporting over-the-top (OTT) services, this extends to having most or all of their channels also available as live streams via cloud and content delivery networks (CDNs), even the largest of live events (such as 2017's Superbowl and the Royal Wedding). In this respect, the cloud can be said to be mature for live delivery. In terms of end-to-end live in the cloud, while this has been possible and in use for more than half a decade, it is still far from commonplace. Software-defined networks (SDNs) and SMPTE ST 2110 and 2022 are accelerating the ability to move to fully virtualized live environments; however, this is one of the areas with the furthest to go in terms of replacing legacy networks and infrastructure with a fully end-to-end cloud equivalent. Conversely, it possibly has the most to gain from replacing very long term and high capital cost infrastructure with targeted operational cost and highly flexible scale.

Post-Production

Every phase of post-production now has some sort of cloud equivalent, and in some areas, the cloud offering is highly mature. VFX is the largest user (by volume) with major studios, productions and facilities now using cloud-compute as either a regular overflow for on-prem render farms or in some case the primary render capability. VFX rendering in the cloud has reached a highly sophisticated phase of deployment where some users will monitor the pricing of available cloud compute in realtime and push render jobs to the most cost-effective cloud at that moment. Since most large public clouds now have bidding marketplaces for unused resource, this can be an incredibly cost-effective approach. It does require highly developed monitoring and bidding systems, and this level of automation points the way forward for the industry to take full advantage of the highly flexible nature of the cloud and the operational advantage of only using (and paying for) the exact amount of processing power you need when you need it.

As long ago as 2015, end to end workflow from camera to screen in the cloud was demonstrated by the University of Southern California's (USC's) Entertainment Technology Center on Abi Damaris-Corbin's film "The Suitcase." Not only did this project set out to prove end to end in the cloud but it did so with stretched goals, including going direct from camera to cloud edge in RAW, the use of high dynamic range (HDR), 360 Virtual Reality (including live from on set), deep metadata including ontology, Academy Color Encoding System (ACES) workflow, the use of Cinema Content Creation Cloud Identifiers (C4ID), and much more. This pointed the way to realistic production workflows of the future which, four years on, are becoming more prevalent. However, it is fair to say that most major functions in post (edit, graphics work, color grading, and sound mixing) are still being performed with local equipment and processing power. While the backend compute power for intensive operations, such as rendering, is moving steadily to the cloud, most realtime work still happens in laptops, workstations, and storage systems that sit inside the facility. This is in part because of the need to amortize existing assets and in part because the complexity of end-to-end still requires much effort to plan and implement compared to known traditional workflows.

Mastering and Versioning

Mastering of content is in a phase of rapid evolution with a steady rise in the adoption of SMPTE-packaged media formats, such as Interoperable Master Format (IMF) and Digital Cinema Packages (DCPs). Alongside this, versioning has been identified as one of the highest growth areas in the media industry, thanks to global markets for content opening up through streaming platforms and an explosion of consumer formats. In this context, automation of versioning and mastering has become a business imperative and new levels of scale that simply were not required in the past are now essential to the process. Cloud solutions in this space include everything from management of localization (foreign language dubbing and subtitling, title graphics, local accessibility tracks, etc.), collaboration on versioning, asset validation, creation of packaged media (IMF and DCP), quality control/quality assurance (QC/QA), business-to-business distribution and direct to platform delivery. In addition to being cloud native, much of this process can be automated, vastly increasing the speed of delivery and dramatically reducing the error rates. In this scenario, the master version of a piece of content is created in the cloud, all versioning happens within the same cloud and the final localized versions are delivered to the end distributor such as a streaming platform again without leaving that cloud. This is an excellent example of an area where cloud technology enables a paradigm shift in business and catalyzes growth.

Delivery

Distribution to the audience is an area in which the cloud is an essential fundamental technology. For streaming platforms, it is not feasible to build global distribution without cloud and CDN. For some in this space, it is their owned network, but this is the preserve of a very few of the largest players, and in most cases, the larger public clouds provide the backbone for streaming services. The biggest challenge here is in "last mile" infrastructure to support the traffic. Currently, more than half of all internet traffic is now delivered over CDN and around 77 Ebytes of video traffic traverses the internet every month. It is predicted that it will grow to as much as 400 Ebytes by 2022. This is certainly a driver for the investment in infrastructure and proponents of 5G mobile networks point to the increase in video traffic as a major justification for next-generation networks allowing significantly improved bandwidth to mobile devices. It should also be noted that this traffic is highly asymmetric, i.e., there is far more download than upload. That being the case, it could be expected that CDNs will grow commensurate with traffic. However, there has been the adoption of large-scale Peer-to-Peer (P2P) networks for content delivery, for instance, in China. This is a low-cost content delivery system (CDS) that could shift the content delivery away from the current reliance on owned CDN infrastructure, but there are still concerns over control and security that are not resolved with P2P technology. Regardless of how content gets there, the cloud will undoubtedly remain in the pivotal hub for large-scale storage and delivery of consumer content over IP for the foreseeable future.

Analytics

Big data analysis is one of the largest uses of cloud across all industries, and it takes advantage of the elastic databases, processing, storage, and AI tools now available natively in many clouds. Coupled with constant connectivity to individuals and windows into their activity, insight from mining these huge datasets is having a huge impact on many aspects of our everyday lives. From retail, news, political campaigns, state policy, and much more, the influence on day to day activities is both enormous and, in many cases, invisible. The media industry has equal amounts of both opportunity and responsibility in this particular area. The technology itself hinges on gathering the user data from services and browsing activity—often linking analytics from multiple sources such as social media, retail, and viewing habits to build a picture of predicted preferences for each individual. This then drives everything from hyperlocal (i.e., the individual viewer) and global decision-making regarding content commissioning, marketing, scriptwriting, casting, and so on. This is largely hinged on AI (deep and machine learning) analysis on huge data sets both structured and unstructured, with highly variable database and processing requirements. In every aspect, this is perfectly suited to the cloud, from the connectivity to scalability and the native services for big data analysis. However, there are challenges to analysis. First, the privacy law, which varies from strong regulation [such as the European General Data Protection Regulation (GDPR)] to almost no regulation in some states, can have widely varying impact on what data can be gathered and how it can be used. Second, and arguably, more importantly in the long term, there is a societal question regarding the way in which algorithms can potentially reinforce biases by making simplistic decisions about how to best win and retain eyeball time from individuals. This is an area in which the technology itself is not necessarily intrinsically harmful but its use could prove to be. However, that is not the subject of this article, but is an area that deserves highlighting and broader awareness.

Challenges of Current Cloud

Even in an ideal media cloud environment, there are still some missing parts when considering the entirety of the content life cycle. Workflows and process are still required to be completed procedurally on a step by step basis. One of the bigger challenges in any media workflow is timely access to source material along with effective communication and productive collaboration between parties and the underlying toolsets they rely upon. The challenges at the front of the mind for those deciding to transition to cloud services are somewhat perennial (see **Fig. 2**) and include how to get data into the cloud, storage costs, and operational challenges, what tooling actually exists and how to leverage automation or people.

Currently, toolsets are typically disconnected. This is due to the fragmented nature of all stages within each phase of a project's production cycle. Each of these phases requires different scopes of work, each within a unique ecosystem containing microprocesses unto themselves. Different departments in filmmaking have such disparate processes because of their different needs, although actually, a large proportion of these unique steps could share substantially the same underlying processes and resources.

In the current paradigm, when a process is completed, it is very hard to retrieve any additional information which was not specified from the beginning. We often hear "I've completed my film!" once shooting has wrapped. However, at this point, there are only raw files which still require a whole host of further work and processing. In post-production, once editorial, VFX, audio mixing, and color grading are completed, there is potentially only one click of a button to "Export to Digital Source Master" and hey presto! The movie is finished again. After QC/QA, the operators say "we are done, now the movie is finished, let's go start another project." In reality, the project is still not yet complete because there is another whole world of processes for distribution, and so it goes on. With the cloud, the closeness of collaboration and the overall content management flow is changing, and we wish to understand how we can improve the complete workflow without burdening it with unnecessary effort. It is difficult to decide what kind of data should be kept and for whom. Of the data that is kept, what should be exposed and when? How should this data be understood or acted upon by storage solutions and how is it governed by an overarching workflow or process. If we have a deep pool of data, then the next step may be to automate processes. If we do not execute some action based on control, we lose

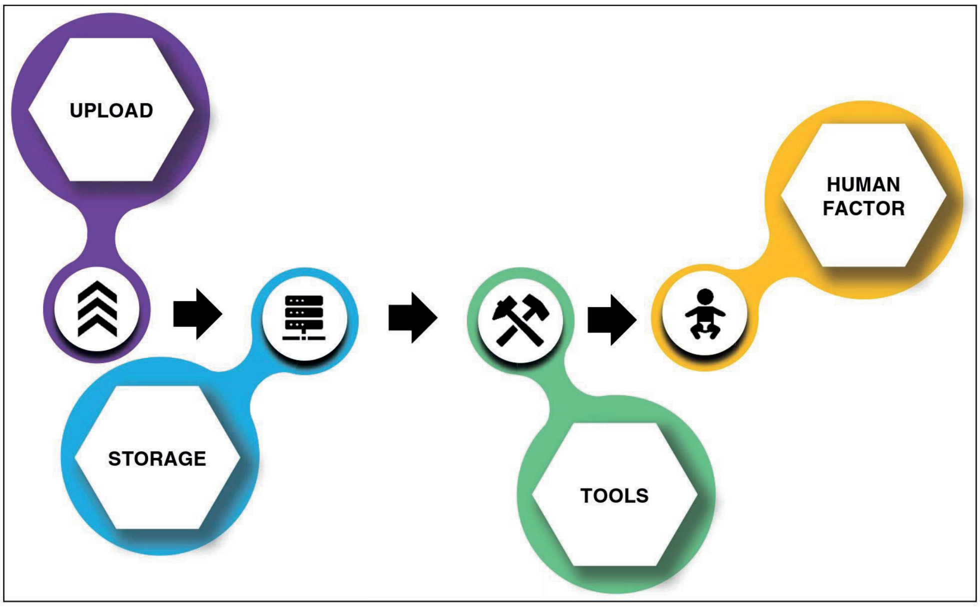

FIGURE 2. The challenges of current cloud.

much of the advantage of the cloud, and it just reverts to a shared planning tool. Therefore, in this way, the cloud presents us with greater possibilities to manage and share content in a user-defined and controlled manner.

Upload

To start any cloud process, source files need to be uploaded. For instance, a typical episodic production generates around 120 Tbytes of data. A feature production may generate more, up to 300 Tbytes. This may not be big by data-center measures, but it is definitely large from an instant availability point of view. To upload 120 Tbytes of files to cloud storage, using Public Internet or P2P, takes time. In an ideal scenario, if the uploader has 1 Gbits/s connectivity, files can be ready to use within 12 days, 5 hours, 12 min, and 11 sec. Tech companies are working to enable this speed on 5G technology, which is hundreds of times faster than 4G; however, we are not there yet. In a remote location, an optimistic and maximum speed is 100 Mbits/s or even lower, so to upload 120 Tbytes of data could take around 122 days. We still have to look at efficient ways for upload to happen, but in most cases, this is still the main reason why the whole post-production process takes place on the ground in post-facilities after post files are uploaded to cloud as backup. In an idealized scenario, we should be taking the camera RAW files directly into the cloud and working from there. Because we are not in an ideal world yet, some solutions include the recently launched Apple Pro-Res RAW, in part looking to address both bandwidth and management of the individual frames in upload to cloud. Certainly, stepping through solutions such as this will bring us closer over time to a true end-to-end cloud workflow.

Storage Cost

Operational Expenditure (OPEX) versus Capital Expenditure (COPEX) for storage. Have it or Buy it?

Often there are questions surrounding the cost of storage and quick access to files on the cloud. Looking at average pricing for top providers 1-Tbyte data costs $23 per 1 Tbyte/month, so for 120 Tbytes, the average cost is $35,000 per year. Of course, after post-production, files can be stored on a long-term basis, where the average cost is $6.5K per year; however, this calculation is only approximate, and cost can vary in different regions around the world. Stakeholders may say "I have already paid for it" and can buy new shiny petabytes of storage every year; however, there are costs involved in risk modeling which includes technical obsolescence, hardware failures, loss of staff, insufficient budget, accidental loss, underestimation of resources or effort, disasters such as fire and flood, and so on.

Tools

Most tools and services, as we know them, are not available as a native cloud solution yet. Even when they do exist natively in the cloud, there are always some missing buttons and functions and various limitations of browsers, security restrictions, and so on. Users cannot work in the same way as they are used to in facilities. For the cloud, files require more preparation in the first place to be able to be used at all. Operators do not work on uncompressed files to preview content anymore, and cloud tools, which use browser-based user interface (UI), have to display proxies. Some people say a proxy can be good enough; however, there is always a technical "but." For example, questions are raised concerning eyeball QC: What conditions are required? Who has access to it and what can I trust? This is why QC of masters in current workflows is being signed off on the "ground" streamed from local storage. This produces additional costs and time stemming from downloading files from cloud to local storage.

Human Factor

The ideal cloud can provide a framework where people can be employed globally and have remote access to projects. There are many things to talk about in this area; however, here we wish to focus on the three main components: skills, collaboration, and procedural change.

Operators skills and knowledge are important to manage any workflow or complete any tasks. Operators can make many unwanted mistakes on all levels and working in the cloud environment—every mistake is visible as it comes with a bill for using microservices. For instance, on the ground, any human error can be run again without visibility of management or the finance team and, thus, with no perceived consequences. In fact, the attrition rate caused by human error may contribute a fairly high percentage to the costs, but without audit access to tools in the media suites, this is very hard to calculate and it is essentially ignored by most businesses (unless the percentage of mistakes becomes so high as to be obvious). Even the waiting time for machines to complete tasks can amount to a huge cost, but again it is very hard to calculate and time-and-motion studies in this area are infrequent at best, nonexistent at worst. A perceived downside of cloud is that all these problems are laid bare in cash cost to the business. It could be argued that these costs were always there and this is an opportunity to reduce or remove them, but that does not do much to dull the pain of the hit to the bottom line.

Collaboration is another important part; we used to talk about our work face to face or in the same room with a group of people. It is well understood that it is easier to come up with creative ideas and solutions for problems with teams that work together in person. As mentioned above, the cloud has no limit to where individuals from the team are located; this has required ongoing learning and improvements in collaborative technology to make the most of how we work and how we collaborate in this connected but disparate future.

Because of a hard stop between the stages, studios have created a new role where it is the remit of technical staff to oversee the whole process from production to distribution and beyond. This continuity and horizontal view of the entire process allows for consistent recycling of information and smoother communication flow. Processes and technology can be seen and evaluated from a different perspective, which, in turn, improves workflows and leads to new ideas, problem-solving, and constant innovation within the content life cycle.

Procedures are changing for security and the way we used to work for the last 100 years in production framework. We already had a transition from tape to digital files and it was successfully navigated, arguably improving the security overall. Now is the time when we move from physical rooms to the virtual environment and much work is now taking place to address how this is assessed and again, where we can improve on the old model. Many audits from local and international security bodies for the media industry now have a very specific focus on cloud and as new workflows are invented, so too are the required security measures that are needed to enable them. For instance, studios are now able to traverse safely from a heavily firewalled traditional facility environment to public cloud without passing through the open internet, even within a virtual private cloud (VPC) and such abilities open up what content can be worked on and at which stage of production.

Big Things That Will Change Cloud

Media Native Cloud

As discussed earlier, there is a standing differentiation between cloud-native applications and forklift or lift and shift cloud usage. However, this is a generic distinction that applies to all industries, as a challenge it is not the preserve of media. However, media does have a very specific challenge with cloud in that the content itself is not understood or natively handled in a cloud environment. A single piece of media in all its forms from RAW, uncompressed intermediates, masters, mezzanines, proxies, deliverables, and localized versions represents a huge amount of complexity in terms of both the understanding of the technical links between these forms and also the contextual links. There are further challenges with the underlying architecture of cloud storage; for instance, the sequential nature of file names for frame ranges of uncompressed material can in itself trip up cloud storage. Cloud storage, while virtual, still relies on physical drives or flashcards for storage, and the placement of files in the physical storage array is linked to the file path [which looks like a uniform resource locator (URL)]. The scattering of files (or objects/blobs as they are commonly described by the main cloud services) across the array looks to the randomness of filenames to spread the files across the storage array. However, if the paths are very similar then the objects will be physically in the same area of the storage array. It becomes obvious at this point that uncompressed frame ranges will cluster in the storage array and thus will have severe limitations when it comes to recalling the objects in one go because a relatively small portion of the storage array is suddenly being saturated with requests and file I/O. Currently, all of these challenges are addressed with the use of asset management software layers, media-specific file systems, and traffic optimization tools that are employed both in the cloud and on-prem. In a media-native cloud, these challenges will be addressed directly by the underlying platform services and become natively available to anyone using that cloud. Beyond being content-aware and optimized for the subtleties of media file structures, this could also support just-in-time (JIT) encoding of main formats to simplify the library, and support an abstracted metadata structure that would simplify (or possibly remove) database requirements. The IMF format potentially has a big role to play here and there have been proposals for databaseless media libraries backed by IMF's highly extensible metadata capabilities.

Automation, Automation, Automation

All industries are moving toward a new level of automation, thanks to the cloud—and media can also take advantage of this. One of the most significant advantages cloud brings to automation is elasticity. In any automation pipeline, the process is limited by the infrastructure [be it a single computer, server farm, or compute shared across a local area network (LAN) or wide area network (WAN)]. This leads to queuing and prioritization of jobs and traffic challenges. In the public cloud, this can be completely dispensed with, instead simply requesting more resource from the cloud as more jobs come in. Depending on the implementation, this can take full advantage of the entire resource of a cloud, although this typically uses more expensive on-demand cloud resource and is limited to applications where service-level agreements (SLAs) are tight, and turnaround is more important than lowest possible cost. A compromise between elasticity and cost uses marketplaces such as Amazon spot instances where the spare resource is available to the highest bidder. This leads to lower costs, but there is always the possibility of the compute resource being taken away at any moment by another customer requesting that resource at the higher on-demand price. The fleet management of compute will define the resilience of the "spot"-based service to fluctuations in available resource and robustness of job completion. With no more reliance on queuing of individual jobs and prioritization, automation can deliver huge scale and greatly reduce turnaround times on traditionally serviced processes. This is likely to lead to a highly disruptive period, where expectations on turnaround time for various services are changed to the point that it becomes difficult or impossible for the manual equivalent to compete. This has already been borne out in the transcode market which is dominated by cloud service providers who have reduced both cost and processing time exponentially.

There are always question marks on job losses but invariably history shows that new technology creates more jobs than it destroys. When Visicalc was launched for Apple, it not only had a huge impact on sales of the computer but also led to the redundancy of hundreds of thousands of accounting clerks. It is estimated that there are 400,000 fewer clerks in the U.S. than there were in the early 1980s as a direct result of the advent of the spreadsheet; yet there are 600,000 more accountants, whose jobs are enabled through the very same software. Of course, with magnified efficiency comes magnified risk, and clearly, automation will just as rapidly propagate errors as it will speed up the processes. Rubbish in, rubbish out is the mantra of automation. What is needed here is intelligence to filter and fix assets upstream of process automation.

Beyond Analytics and Image Recognition

Currently, AI is being widely employed by the media industry in the cloud, and, as previously discussed, this is a space in which cloud is the natural home for these services. However, there are numerous applications being developed in the creative space that employ AI to perform tasks thought the preserve of humans. We will see an explosion in the use of AI to work as creative assistants in all areas including edit, color, VFX, sound mixing and more. Furthermore, AI is starting to play a significant role in improving processing the overall quality of the content beyond what can be achieved by fixed algorithms, and applications for next-generation image scaling, frame-rate conversion, and compression optimization are already becoming widely available. Of course, these applications could equally be employed locally, but even then they will most likely be backed by a cloud-based training system for scale and reach. Another important area for AI is ensuring the quality and validity of assets in automation pipelines. This is currently still a major bottleneck when it requires human interaction, but that bottleneck can only realistically be cleared through yet more automation. In this case, building AIs to solve the rapid validation of assets, whilst a niche area in itself, will likely have a big impact on releasing the potential of cloud for media processing.

As with any machine learning, the success of any of these applications will heavily rely on the quality of the training data. This not only goes beyond simply having access to content to train on but also on how disparate data sets can be normalized to get the best performance from the trained AI. Here, curation is a major challenge. The larger issue, however, is the availability of the source data itself. Clearly, a well-trained AI has commercial value, and if the quality of the training data is key to this, then, by extension, some of the commercial value of the AI comes from the training data itself. That being the case, existing owners of content are often reluctant to provide these training data sets without some sort of arrangement regarding the intellectual property and possibly commercial agreements regarding monetization of tools. This will undoubtedly lead to a market for vanilla tools that can then be trained to excellence by content owners for their own use, possibly biasing the tools to the content's own style. This will likely be seen as an advantage for any large content owners as they will retain the commercial value of their IP and potentially customized to their overarching look and feel (if it is a specific content house or even individual production). However, this may keep some of the best performing tools out of reach of lower budget or lower volume productions, which might benefit most from the cost savings that might be afforded. Time will tell but it is certain that one of the fastest-growing spaces in both cloud and media is the use of AI and it will be impossible to ignore the impact this will have in the coming years and decades.

About the Authors

Tomasz Witkowski has more than 21 years of experience in the global motion picture industry. Having held positions at Dolby, Éclair, Visual Data, and currently Sundog Media Toolkit, Witkowski is an accomplished engineer with knowledge spanning across post-production workflows and processes, encoding, and conforming and editing theatrical and TV content for mastering. He is a member of influential media bodies such as SMPTE, the British Standards Institution (BSI), the International Organization for Standardization (ISO), and the Digital Production Partnership (DPP). Witkowski was a key part of the team that set up the SMPTE Poland Section and helped create the portal for SMPTE DCP. He held a migration project and co-founded the Soho Media Club. He is involved in many technical projects, such as IMF.

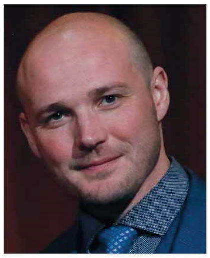

Richard Welsh is the co-founder and the chief executive officer of Sundog Media Toolkit Ltd, U.K. Based in the U.K., Welsh has been in the cinema industry since 1999. He has been involved in various technology development works during this time, primarily in digital cinema, and is named on patents in the area of visual perception of 3D. He set up Sundog in 2013 and, along with his colleague, Christian Ralph, created the Sundog cloud platform. Sundog specializes in versioning automation, hyperscale processing, and AI applications for media. Sundog works with several major Hollywood studios and broadcasters and contributed to feature films, broadcast series, documentaries, and supporting student films from universities around the world. Welsh has previously held positions as the head of digital cinema operations at Technicolor, U.K., and the director of digital cinema at Dolby Laboratories, U.K., and he has served on the board of SMPTE. Welsh holds a BSc degree (Hons.) in media technology and an honorary doctor of technology degree from Southampton Solent University, U.K.