

## Video Production in the Cloud with EditShare and Amazon Prime Video

### Now is the Time to Tell Stories in the Cloud

Since the days of the first cave people, humans have been sharing stories through visual mediums. With the advent of photography, film, and later television, storytellers could communicate their stories to the masses. And now, in the last decade we have seen fundamental transformations in the tools people use to tell their stories, as well as in the ways they can convey them - enabled by advances in technology such as mobile devices with built in high quality camera capabilities, drones available at consumer price points, high resolution video formats, inexpensive video and audio editing applications, advancements in digital effects, and new distribution networks to reach audiences around the world. As a result, we have seen an exponential growth in the number of stories being told and also in the audience that can be reached for any given story.

Obviously, the cloud has played a big role here. To reach audiences in more flexible ways we have

seen significant adoption of cloud technologies in the video delivery chain with encoding, packaging, encryption, distribution and many control functions deployed in public and private clouds. These solutions take advantage of the massive scalability of the cloud and advanced features such as artificial intelligence and machine learning.

Visual effects and animation workflows have also greatly benefitted from the massive rendering power of the cloud with large farms of resource being spun up in minutes and torn down when no longer required.

The production, editing and post-production side of video, however, has not really made as significant steps into the cloud. A number of factors have frustrated desires to build a full video production workflow in a fully cloud native environment, not least the transit, security and storage of high resolution raw media, as well as the perceived cost.

The benefits of cloud for production workflows are tantalizing. With many productions having distinct workflow needs, the option to spin up an environment to fit the requirements of a particular project and then reconfiguring it for the next project would be a game changer for many organizations. In addition, the short-term nature of many projects – think of a TV series pilot or one-time event – add to the desirability of being able to quickly turn on and then turn off the whole production environment without having to make any capital investment. Collaboration can also be enhanced with the help of the cloud by allowing remote workers to participate in all aspects of content creation, including editing, from anywhere in the world. Finally, the opportunities to collaborate are also magnified when the barriers of traditional IT are broken down by an always accessible, always secure system that is available where you need it, when you need it.

As we look across the wide range of verticals that produce video, we continue to find opportunities to benefit from an on-demand, flexible, workflow-centric approach enabled by the cloud. Be this the opportunity to aggregate raw news clips from across the globe, to enabling sports editors to generate highlights from the comfort of their own homes, education institutions allowing students

access to the projects and reality TV shows taking advantage of advanced AI and machine learning to wade through growing volumes of raw footage.

### Create Content in the Cloud

With our heads in the clouds, dreaming of this advanced video production environment, we must first have a quick reality check: is it really possible to create content remotely in a cloud deployed system? And even if it is possible, are the hypothetical benefits what they appear to be?

To help us answer those questions, EditShare partnered with Amazon Web Services (AWS) and Amazon Prime Video to create an end-to-end remote production environment for the generation of highlight packages for the ATP Monte Carlo Masters tennis tournament.

Beyond any technical success of deploying this solution in the cloud, a key evaluation consideration was whether we could solve a real world challenge and demonstrate additional value and/or cost savings that the flexibility of cloud could deliver to the Amazon Prime Video sports production team.

Expertly created highlights packages provide significant value to the economics of sports TV. With live sport being streamed to customers



IMG PRODUCTION CENTER IN LONDON DURING THE ATP MONTE CARLOS MASTERS TENNIS TOURNAMENT





*EDITOR WORKS REMOTELY FROM LOS ANGELES*

over the Internet, highlights packages provide both interstitial content for use within live commentary and additional bite-sized packages for consumption on second screen and alternative consumption destinations. The demand for this type of content continues to grow, but the traditional monetary and environmental costs to fly video editors and producers to sports venues becomes a significant barrier to the ongoing viability.

Ready to think differently about the problem, Amazon Prime Video approached EditShare to explore the potential of deploying a solution in AWS that would allow them to ingest live tennis feeds into the cloud, storing the high resolution media in the EditShare File System (EFS), while ingesting the content into the Flow Media Management system for indexing, meta tagging, content approval and finally archival.

The aim was to allow producers, editors and supervisors to work from anywhere in the world, all accessing a centrally managed solution that was both high performance and cost effective.

The cloud workflow started with live video sent directly from the venue over fiber optic cable to the IMG Production center in London. While

we could have uploaded the live switched feeds directly from Monte Carlo to the cloud, it was easier in this case to intercept the video using the existing channel back to London.

EditShare Flow ingest servers captured the live stream and uploaded the video using UDP file acceleration into the AWS data center in Virginia.

The captured video was then stored in virtualized EditShare EFS storage servers and indexed using Flow. Editors in Los Angeles were then able to connect via remote desktop into cloud workstations running Windows and hosting Adobe Premiere Pro and EditShare management software to connect to the virtualized high performance storage and media management software. To reiterate, the whole production environment was in the cloud, and the editors were simply connecting in via remote desktop software.

With an experience indistinguishable from using the non-linear editing software on their own desktop machines, the editors crafted high quality highlights that were rendered directly into the EditShare cloud storage and then transcoded automatically into a review-and-approval-friendly

low-resolution proxy version. The proxies were then viewed by directors in London using the online AirFlow tool, and any needed adjustments were noted on the clip for reworking back in Los Angeles. Once approved, a Flow automation job would deliver the high resolution edit back to the playout center in London for inclusion in a live show.

The benefits of this workflow were clear to Amazon Prime Video: the editors were producing timely, high production value highlights values that added to the enjoyment of viewers around the world, without costly travel and other overheads.

The benefits to EditShare were also clear, because in one multi-faceted project we were able to evaluate many different building blocks of cloud production and workflow that might ultimately be combined in different configurations to create workflows for news, reality TV, advertising, film and much more. Not every workflow will require live ingest, some workflows may work well with remote proxy-based editing and some may seek to build deep archives in cloud object storage. Key to enabling this is flexible cloud video software that, if implemented well, can truly simplify the production chain for the next wave of storytellers.

## Video Production Technology is Cloud Ready

There are a number of key technology enablers that pave the way for this transformation, from high bandwidth remote desktop technology to software defined storage solutions, from network streaming protocols to microservices for automation and asset management.

A key foundation is flexible storage that can provide high streaming throughput for editors to work with vast quantities of high definition video clips as if they were local to their workstation. This storage must not be limited to the capacity of a single machine instance and cannot rely on hardware technologies such as RAID cards. It must be scalable and hardware independent.

The EditShare File System provides a software platform to abstract the underlying hardware or

virtualization platform and has been designed from the ground up to flexibly scale out from tens of terabytes to petabytes of storage. It provides a single namespace that spans an entire cluster of servers that contribute to the throughput and resilience needs of video workloads.

The single storage namespace can be carved into unlimited flexible “Media Spaces,” each of which can act as buckets for storage of media for a particular project or group of users. During the lifetime of a Media Space, the storage goals can be modified, seamlessly moving assets from spinning disks to SSDs when required. In the AWS cloud, the EditShare File System can be utilized to build clusters of nearline and online storage from the available storage types, for example throughput optimized spinning disk (e.g., AWS EBS st2), SSDs (e.g., AWS EBS gp2/io1) or instance attached storage.

The ability to leverage a wide variety of storage types provides the flexibility to trade performance against cost, with the clustered file system always providing an abstraction to allow migrating between storage types as workloads change. A very cost effective solution with good throughput can be achieved with low cost block storage and a larger number of low cost EC2 instances, benefiting from the combined throughput of storage spread across the network file system. To date, in EditShare deployments within AWS, throughputs of up to 2000 MB/sec to each virtualized workstation have been achieved, proving this architecture can support anything from proxy editing to real time uncompressed 4K finishing work.

Understanding the dynamics of different EC2 instances and storage types is key to choosing the best underlying AWS configuration for a given workload. For example, AWS provides options to optimize throughput and latency to its block storage (EBS). In addition, each different instance type and size provides a different level and balance of guaranteed and burst networking from the instance itself. Clustered storage for video applications typically requires little CPU, only a modest amount of memory but mostly sustained

network throughput (offered through enhanced networking support of certain instance types).

With a solid scalable storage foundation, one of the next key challenges is transporting high resolution media from a source location into the cloud. While it has become commonplace to transit MPEG encoded video for live broadcast streaming applications, such streams typically run at data rates under 25 Mbits/second. This is compared to a SMPTE 292M HD-SDI input stream with a bit-rate of 1.485 Gbit/s, and to compressed ProRes, DNxHR and XAVC streams that typically run between 100 and 440 Mbits/sec.

For the ATP tournament use case, EditShare used a Flow Ingest server to capture 4 raw SDI feeds from the IMG video production facility in London and encode them into 5 minute chunks in XDCAM 50 Mbit/s format (50 Mbit/s was chosen as a compromise between good origination quality and a bit rate that wouldn't overwhelm the available upstream network connectivity from the facility). Once each chunk of the recording was completed, a Flow Automation job would upload the resulting file to the EditShare EFS storage system running in AWS. ExpeDat Fast File Transfer technology was used to accelerate delivery of these files across a standard internet connection. ExpeDat was chosen in this case, however many other file accelerators such as Aspera and Signiant provide similar capabilities.

The individual file chunks became available immediately for editors to work with on their virtualized Windows workstations (accessed by Remote Desktop). In addition, the chunks were scanned and ingested into the Flow Media Management solution hosted in AWS, allowing video editors to tag clips and rapidly assemble highlights. While there were some drawbacks to having multiple files per hour of footage, the Flow solution provides capabilities to easily link clips together before editing.

As an alternative to file chunking, there are also a number of streaming protocols such as Secure Reliable Transport (SRT), Reliable Internet Stream Transport (RIST) and solutions such as Zixi that

provide quality of service offerings plus advanced end-to-end stream monitoring capabilities.

Depending on the throughput/latency requirements, many public clouds provide dedicated leased lines, such as AWS DirectConnect. While this was not required for the ATP use case, it can provide significant benefit if low latency and dedicated bandwidth are required.

For a large video archive, devices such as an AWS Snowball can be used with a local file transfer onto a device that is then shipped directly to Amazon and ingested into S3 object storage. From there it can be scanned by Flow for searching and retrieval into online storage for future editing.

The final piece of the jigsaw is to provide ultra low latency remote desktop access to video editing workstations hosted in the cloud. While it is certainly possible to perform a proxy edit workflow, where the remote editor streams low resolution versions of clips for editing on a local workstation, the opportunity to work directly with the high resolution media in the cloud is highly appealing.

The PC-over-IP (PCoIP) protocol, provided by Teradici, makes this a reality. The editors from Amazon Prime Video needed only an inexpensive Mac or PC and approximately a 20 Mbit/s connection to the cloud editing environment, and they were amazed by the quality of the experience and were able to use all the same tools as they do today on their local workstations. In addition, the benefits of being able to manage multiple workstations from a common template makes onboarding of new editing staff into a project straightforward. In addition to using a common VM image (e.g., an AMI), cloud providers such as Amazon provide a full managed virtual workspace solution that incorporates management capabilities, single sign-on desktop access, and much more.

There is one more important thing to note about the EditShare / Amazon Prime Video cloud project and that relates to AWS costs. In

specifying the EC2 instance and storage types to be used in this project, EditShare determined it was possible to get the required performance for multi-user multi-layer timeline editing mostly using the very smallest r5 instance types at an operating cost of 12.5 cents/hour per instance plus the monthly charge for 48 TB of Elastic Block Storage (EBS). With 5 x EC2 instances (one being slightly more powerful for Flow) deployed overall, the total price to operate the system 24/7 came down to about \$94 per day, plus the cost of the Windows graphics workstations (in this case, overspec'd at \$27/day per workstation) and a small amount of egress charges for the Teradici Remote Desktop access. These prices do not include EditShare EFS and Flow software, nor do they include NLE software. The prices are for AWS infrastructure only and were from the AWS North-East/Virginia data center in May, 2019, and reflect on-demand pricing without any discounts that can be negotiated based on multi-month utilization commitments. When compared to the total cost of an on-premise deployment that requires rack space in an air-conditioned server

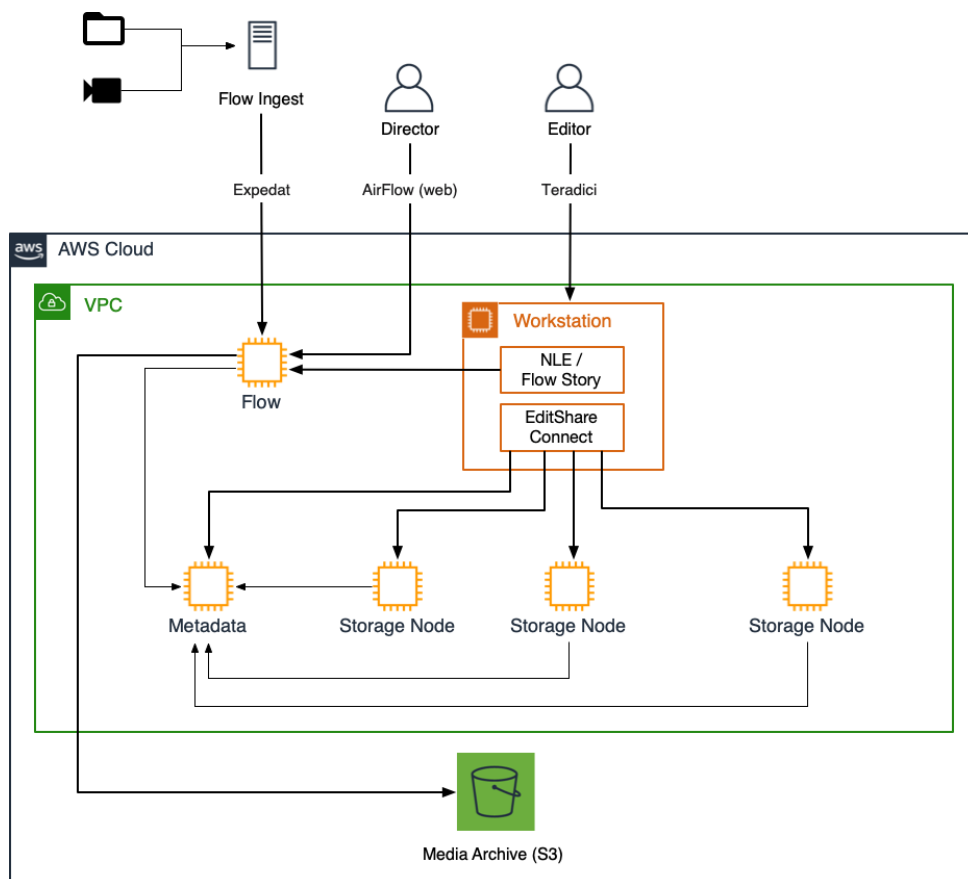
room, electricity to run the servers, physical edit suites, building wiring, common space for each person who works in the building, IT personnel to maintain equipment, etc., the cloud deployment is a bargain. And cloud costs look as if they will be coming down in the near future as new EC2 instance types are introduced that offer a better balance of network connectivity, disk storage, CPU and RAM that are more optimal for the video editing use case.

### A Brighter Future in the Cloud?

Having established in the EditShare/Amazon Prime Video project the potential to run video production in the cloud and the technologies that enable it, we can now consider the factors that will make a step change improvement in the storytelling workflow in the cloud.

#### Cloud Flexibility

A key benefit of true cloud solutions is flexibility – in specific, the ability to craft an environment to fit the specific demands of a particular workflow and to adapt to the needs of a particular production



A HIGH-LEVEL DEPLOYMENT ARCHITECTURE FOR END-TO-END WORKFLOWS

or project. Achieving this flexibility requires a full automated deployment that can spin up resources when needed and shut them down when idle or no longer required. If you have a sudden requirement to restore a large amount of content from an archive to online storage or add 20 new video editors to your production, a cloud native solution can adapt on demand. The need to plan the full details and scaling of your system carefully upfront are much diminished, now you can even fix that in post!

### Open APIs

To realize this benefit, an open approach is essential. While it is desirable to provide a validated reference blueprint for a cloud deployment, each customer will want to tailor their environment to fit their needs. Each vendor in a cloud ecosystem needs to embrace Open APIs that allow customers, integrators and other vendors to easily pull together a solution that fits the needs of sophisticated workflows.

Technologies such as the OpenAPI Specification (previously known as the Swagger Specification) enable great interoperability with APIs documented in a machine interpretable format that can be used to generate stub implementations and SDKs in multiple programming languages (as well as generating interactive API documentation).

### Overhead Management

With each of the cloud vendors, you can find infrastructure and platform components on offer that help reduce the overhead of managing a cloud production solution. From fully managed databases to machine learning and artificial intelligence services, the options are vast. This presents both opportunities and challenges.

For instance, removing the overhead of managing day-to-day technology operations - eliminating the need to worry about upgrades, security patching and back-ups - is very appealing but there currently isn't universal support across cloud vendors for many of these services. Most cloud providers have some form of SQL database as a service (usually MySQL but often Postgres, Oracle and SQLServer). NoSQL databases have less universal availability, yet many modern applications have been built with a document

style or key/value datastore. The choice then is to decide to utilize a database that is only available in a single cloud or maintain portability between clouds. Often customers and vendors can find themselves locked into a particular cloud vendor, losing the flexibility to utilize new offerings in other clouds or take advantages of price reductions from a particular cloud.

### AI and Machine Learning

An opportunity afforded by cloud deployment is the massive reduction it brings to deploying large scale machine learning and artificial intelligence solutions. These can bring value throughout the production workflow: speech to text can bootstrap clip metadata logging and more sophisticated AI can detect sentiment and emotions, all useful for identifying the right clips to build an initial edit or find alternatives when a particular angle isn't working. Object and person detection can also be used to help quickly find the clips and even support an audit trail for usage rights. There are great opportunities to utilize machine learning, trained on real editing decisions, to provide initial assembly options that can be used by editors as a starting point to craft the right story.

### Security

Unlocking all these benefits is clearly very appealing but security of the content is critical. An advantage of most public cloud environments is there are significant built in security options



AWS PRODUCT ICONS INCLUDE EXTENDED FEATURES IN THE CLOUD SUCH AS AI AND ML



that come at considerable cost and complexity to implement in an on-premise network. From encryption at rest for archive and online storage, to fine grained network level access control, from Web Application Firewalls (WAF) to advanced intrusion detection, from hardened operating system images to DDOS protection. These features have been developed to meet the strongest of security requirements and proven out for many years in video distribution systems. Following the best practices and getting a full review (e.g., an AWS Well-Architected Review and Penetration Test), will ensure that content is safer in the cloud than anywhere else.

### Your On-Ramp to the Cloud

So, to summarize, the benefits of creating content in the cloud are compelling, the underlying technologies required (automated deployment, cloud database, service and upgrade management, AI, security, etc.,) are not only available but are probably more robust than on-premise alternatives, the cost is not unreasonable, and thus the stars are aligning to make cloud production practical today.

That said, for those who are skeptical about simply flipping the switch and moving everything to the cloud from one day to the next, there are hybrid approaches that can support a more phased transition. For instance, the same technologies used in the EditShare/Amazon Prime Video project can be deployed in the cloud solely for disaster recovery/failover purposes while conventional hardware-based EditShare EFS storage and Flow media management tools are deployed on-premise; in this use case, synchronization tools copy media and metadata on a regular basis to a “pilot light” cloud deployment that is spec’d for minimal performance and expense but that can be transformed in a matter of a few minutes into a high-performance system -- by rebooting into more powerful EC2 instance types. The cloud environment is only used actively for production when the primarily on-premise environment is unavailable. In other words, we’re sitting at the on ramp to the cloud and only get on the highway when actually needed.

An even more economical DR scenario is to synchronize content and metadata from the on-premise system to economical object storage in the cloud – for instance, S3 or Glacier in AWS. In this case, the cloud production environment with high-performance EditShare storage is only spun up if a DR situation calls for it, and then Flow restores the required content from object storage to online storage.

For those who wish to go beyond DR but who still don’t want to move high-resolution content to the cloud – because of the time required to transfer high-resolution content or the expense of storing so much data there – another hybrid solution that can be deployed today is to keep high-resolution content on-premise on EditShare Storage and only synchronize low-resolution proxy files and metadata to the cloud. In this scenario, workers situated anywhere in the world can connect to the cloud and enjoy scalable streaming proxy timeline editing with the Flow Story non-linear editor, as well as remote viewing and logging with the AirFlow web application, and then when a story is finished the timeline is conformed at the on-premise location that has the high-resolution media (and if high-resolution content is dispersed at multiple on-premise locations, the required content is fetched from wherever it resides).

Finally, for those who are worried about things like “there’s no way to get an SDI output from the cloud for accurate color grading” or “how to do a Dolby 5.1 mix when Remote Desktops only support 2-channel stereo sound”, the good news is that EditShare EFS storage can be deployed in a virtualized cloud environment that nonetheless allows remote mounting over the Internet with surprising levels of performance. So, in a scenario in which most editors work via Remote Desktop to the cloud, a small number of on-premise workstations that have SDI and surround sound outputs can be put into the mix for the tasks that can’t be performed in the cloud itself.

In conclusion, while content creators have been dabbling with cloud production over the past few years, now is the time to begin the transition in earnest. There are many different scenarios where this makes sense today, and the reasons to hold back are swiftly evaporating.