

Ready, Set, Finish! A Studio and Production Workflow in the Cloud



AWS and AWS Partners Demonstrate an Interconnected Post-Production Environment and Modern Toolset Operating in the Cloud



16

Integrated Partners

10-BIT

End-to-End Cloud Workflow

REAL-TIME

In-Camera VFX

SUMMARY

At the International Broadcasting Convention (IBC) 2022, Amazon Web Services (AWS) set out to demonstrate the progress made in moving professional media production workflows to the cloud by creating and posting content live in the event trade show booth, using real tools, workflows, and participants. The end-to-end production workflow in the cloud was modeled on the MovieLabs 2030 Vision[®] and showcased several of the key components of a real-world production: pre-visualization, shoot, ingest, editorial, visual effects (VFX), compositing, color/finishing, and quality control. The cloud studio environment spun up for the production included deployment orchestration, shared storage, asset management, and workflow automation. This AWS Studio in the Cloud case study addresses some of the most challenging post-production workflows with very large professional media files and provides a roadmap for media companies on how to deploy workflows at scale to the cloud.



BACKGROUND

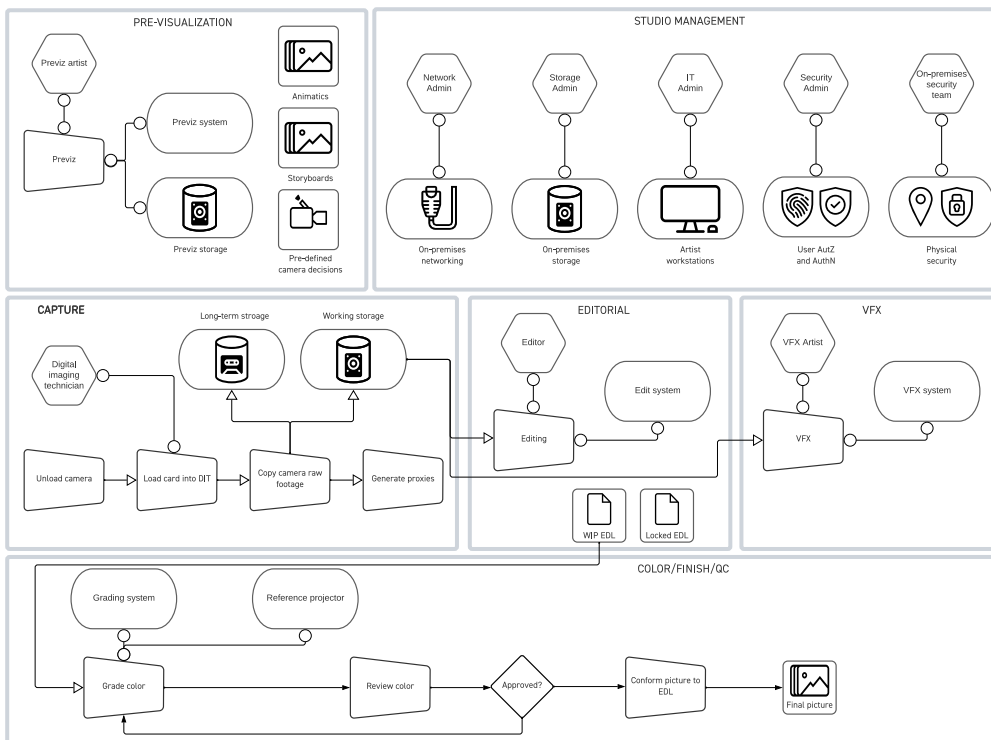
In post-production workflows, physically moving content between environments requires a large effort. This movement creates operational and labor costs, increases the chances of lost files or metadata, and affects the overall security of the environment. The time and resources spent on moving and managing data can detract from adding value to content through creative iteration.

Traditional workflows include manual media movement from on-set to a centralized storage location, which is usually at a post-production facility. Media is copied from camera cards onto physical hard drives and is often manually moved to the production facility for loading into central storage. This causes delays in dailies and editorial workflows because it takes time to move media between locations.

In addition, Quality Control (QC) is manual in traditional workflows. An operator views media in high-quality file formats, which requires media to be transcoded and mounted to be viewed on a reference monitor for QC engineer review. At this point, if there are issues with the media and QC fails, the process must start again, delaying QC completion and the next steps in the workflow or delivery of the content.

The media files are large and often represented by lower bitrate proxies, which introduces new challenges to maintain the connection between the original and related proxies. The Original Camera Files (OCFs) are in a variety of non-standard file formats – often frame-based codecs, which present unique challenges as they can't be viewed easily in a browser or file system.

A TYPICAL LEGACY ON-PREMISES PRODUCTION ENVIRONMENT



Today, the cloud is widely used for storage of media files (especially for archive), file transfers, lightweight proxy editorial, and for large compute jobs like VFX rendering. However, AWS has committed to implementing the Principles of the MovieLabs 2030 Vision and is focused on filling today's gaps to create more effective, efficient cloud-based workflows and demonstrating today's tools with the cloud providing work-in-progress (WIP) storage throughout.

Figure 1: A traditional workflow for a narrative TV show or movie with multiple stages of data movement between departments

SOLUTION

In the last few years, new cloud-enabled tools, open APIs, and the ISV partner ecosystem have expanded the art of the real beyond traditional cloud workloads and towards a holistic production workflow with shared storage. To showcase this readiness, AWS for Media & Entertainment worked closely with middleware and creative application partners to showcase an entire production workflow in the cloud at IBC 2022.

The AWS Studio in the Cloud solution included commonly used media applications (Adobe Premiere Pro, Baselight FilmLight) all running on virtual machines, processing stacks for transcoding and file preparation (e.g., Transkoder), and connections to SaaS platforms (e.g., Moxion for real-time dailies). The entire pipeline ran at 10-bits with color accurate monitors for QC as if the final mastering was happening locally within a production facility connected to a Storage Area Network (SAN). The assets captured were of a live scene with real-time compositing of a background plate from a blue screen. The solution included realistic monitoring, proxies, previews, review/approvals, QC, and transcoding into file formats commonly used in movie and TV production workflows.

AWS AND AWS PARTNERS RE-ARCHITECTED CLOUD-BASED PRODUCTION ENVIRONMENT

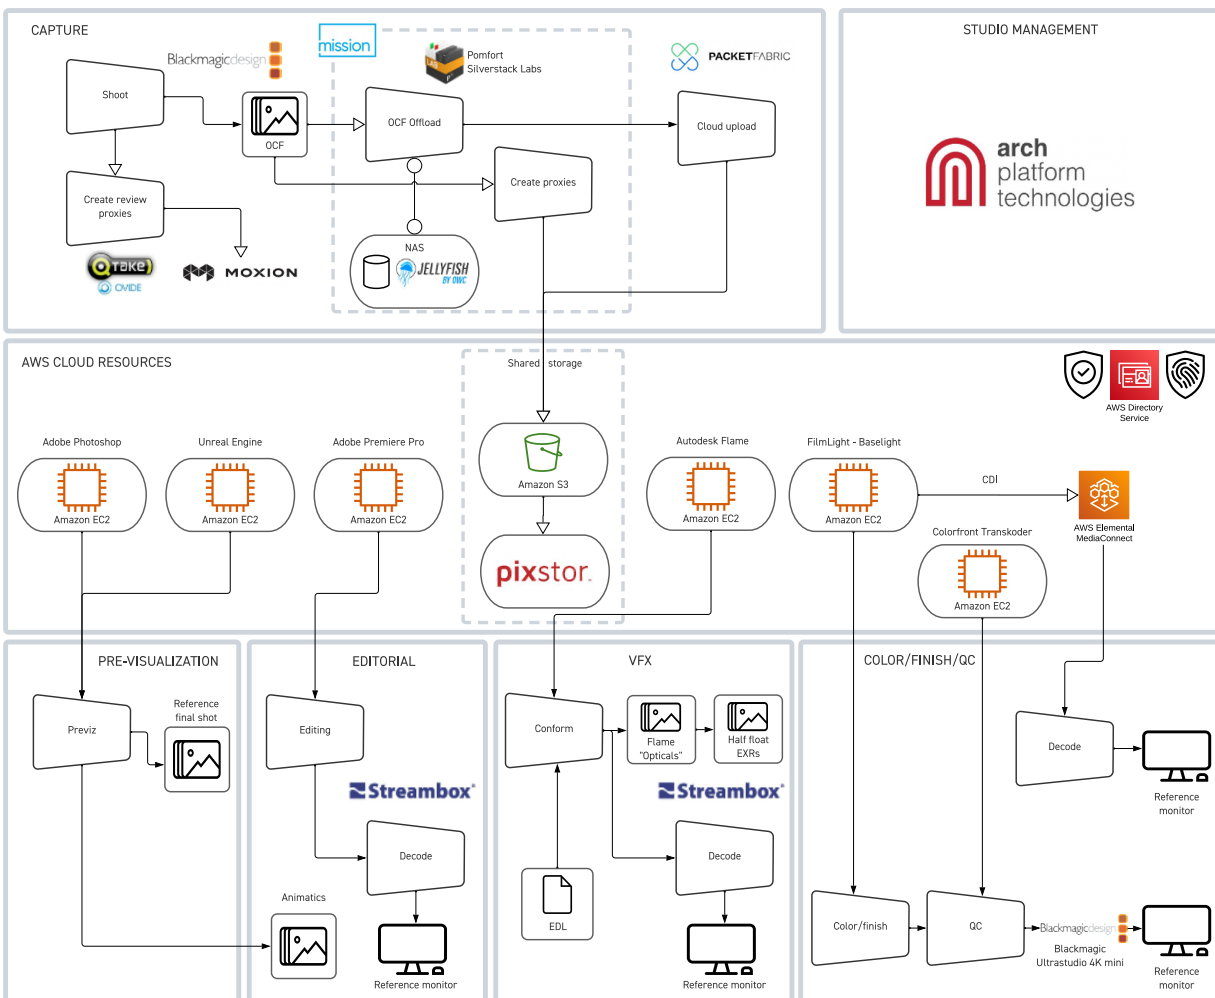


Figure 2: The cloud-enabled infrastructure and AWS Partners used to demonstrate end-to-end post-production in the cloud

ARCHITECTURE

The on-set camera feed was sent to QTAKE which created real-time proxies and automatically uploaded them with accompanying camera metadata to Moxion for review. QTAKE displayed assets and feeds in real-time with on-set compositing for a proxy view of the finished product.



Figure 3: On-set at IBC with Cludia, the robotic dog and star of the show

Meanwhile, an automated process, using Pomfort Silverstack, offloaded OCFs from the camera media and sent them to an on-site Jellyfish Network Attached Storage (NAS) that synchronized to an Amazon Simple Storage Service (S3) bucket for disaster recovery protection. Silverstack generated a proxy file and ingested the files and associated camera metadata directly to Amazon S3 so the WiP copies could be edited. Upon landing in Amazon S3, a serverless compute function —using AWS Lambda— ran automatically to hydrate the media from S3 to the PixStor filesystem. The PixStor storage was powered by three compute-optimized, clustered servers on Amazon Elastic Compute Cloud (Amazon EC2), backed by Amazon Elastic Block Storage (EBS) volumes for roughly 27 TB of storage.

For editorial, the editor remotely connected to a graphics-powered workstation running Adobe Premiere Pro in AWS. The workstation was connected to the PixStor storage cluster over Server Message Block (SMB) protocol that provided low-latency access to the media for the editor to create a rough cut and VFX pulls. Network Device Interface (NDI)¹ was used to transport frame-accurate, ultra-low latency video over Internet Protocol (IP) without compromising quality. In this solution, Adobe Premiere used NDI to send a live view of video playback to a Streambox plugin that sent the video to an on-premises Streambox decoder co-located with the editors, where it was displayed on a high-fidelity, broadcast reference monitor.

¹ Network Device Interface (NDI) is a royalty-free standard developed to facilitate near real-time and ultra-low latency video communication over internet protocol (IP) in a frame-accurate manner without compromising quality.

The Arch Platform manifests key components of the MovieLabs 2030 Vision today. At IBC, our studio management platform underpinned the end-to-end production workflow from content ingestion to distribution, supporting collaborative workloads for video ingest, editorial, VFX, color correction, and review. Being able to quickly build secure cloud facilities for a multitude of content creation workflows is a big boon for remote production, allowing teams to work faster, smarter, and more efficiently than ever before from anywhere in the world.”

LAURA TEODOSIO, CHIEF EXECUTIVE OFFICER, ARCH PLATFORM TECHNOLOGIES

Unreal Engine 5 ran on the latest generation NVIDIA GPU-based instance equipped with four NVIDIA A10 GPUs that handled real-time rendering and compositing of VFX on final plates as well as pre-visualization. Autodesk Flame ran on a slightly smaller graphics-powered instance, accessing the PixStor cluster via the native Linux client. This allowed the VFX artist to collaborate with the editor to achieve a final edit. Interchange files such as Edit Decision Lists (EDLs) and Audio Interchange Files (AIFs) were stored in Moxion (backed by Amazon S3), allowing artists to access them as needed. Flame used the [AWS Cloud Digital Interface \(CDI\)](#)² to transport an uncompressed live video feed to Streambox, where an on-premises decoder displayed the high-resolution image on a broadcast monitor.

When the edit was complete, the final timeline was assembled in Baselight for color grading. Baselight ran on a graphics-powered instance, using CDI to send uncompressed video to AWS Elemental MediaConnect. MediaConnect applied a low-latency JPEG-XS 10:1 compression to the stream, enabling the artist to view the graded image in real-time from the cloud. In addition to the client machine, an IP JPEG-XS decoder box, a broadcast monitor, and a grading panel are needed for the artist to use this workflow. AWS Elemental MediaConnect now supports RGB 10- and 12-bit 4:4:4 color spaces via CDI, which enables workloads where visually lossless fidelity, accurate color, and low latency are required.

Once the grade was complete, Colorfront accessed the content for QC. Colorfront created an encoded stream for the streaming player running locally on a MacMini and connected to a Blackmagic Ultrastudio 4k Mini for output to a reference monitor. This system provided a color-accurate image on-site for QC processes and review and approval without ever moving the content. Each workstation in this pipeline used an NDI screen-capture feed into Moxion, enabling a remote director or producer to have a real-time view of the entire workflow at any time.

² AWS Cloud Digital Interface (CDI) is a network technology that allows you to transport high-quality uncompressed video inside AWS, with high reliability and network latency as low as 8 milliseconds. AWS CDI supports uncompressed video up to Ultra-High Definition (UHD) 4K resolution at 60 frames per second.

Flame ran on a slightly smaller graphics-powered instance, accessing the PixStor cluster via the native Linux client. This allowed the VFX artist to collaborate with the editor to achieve a final edit. Interchange files such as Edit Decision Lists (EDLs) and Audio Interchange Files (AIFs) were stored in Moxion (backed by Amazon S3), allowing artists to access them as needed. Flame used the [AWS Cloud Digital Interface \(CDI\)](#)² to transport an uncompressed live video feed to Streambox, where an on-premises decoder displayed the high-resolution image on a broadcast monitor.

When the edit was complete, the final timeline was assembled in Baselight for color grading. Baselight ran on a graphics-powered instance, using CDI to send uncompressed video to AWS Elemental MediaConnect. MediaConnect applied a low-latency JPEG-XS 10:1 compression to the stream, enabling the artist to view the graded image in real-time from the cloud. In addition to the client machine, an IP JPEG-XS decoder box, a broadcast monitor, and a grading panel are needed for the artist to use this workflow. AWS Elemental MediaConnect now supports RGB 10- and 12-bit 4:4:4 color spaces via CDI, which enables workloads where visually lossless fidelity, accurate color, and low latency are required.

Once the grade was complete, Colorfront accessed the content for QC. Colorfront created an encoded stream for the streaming player running locally on a MacMini and connected to a Blackmagic Ultrastudio 4k Mini for output to a reference monitor. This system provided a color-accurate image on-site for QC processes and review and approval without ever moving the content. Each workstation in this pipeline used an NDI screen-capture feed into Moxion, enabling a remote director or producer to have a real-time view of the entire workflow at any time.

PARTNERS

The AWS Studio in the Cloud would not have been possible without the support and technology solutions from various AWS Partners contributing and interconnecting their systems to demonstrate what is achievable. For the film production component, AWS worked with Blackmagic Design, OWC, IN2Core, Mission, and PacketFabric. The solution used Blackmagic cameras connected to a Digital Imaging Technician (DIT) cart where H.264 proxies were generated using QTAKE and sent to Moxion. OCFs were offloaded to a high-throughput storage NAS powered by Jellyfish (OWC). The OCFs were ingested into Amazon S3 using PacketFabric for on-set stage connectivity to the cloud.

For cloud studio management, AWS worked with Moxion, Arch Platform Technologies, HP, and Pixitmedia. Arch Platform Technologies provided workstation management capabilities for provisioning and managing the lifecycle of the artist workstations in the studio environment. HP enabled PCoIP desktop streaming using its HP Anyware solution that allowed the artists to remotely connect to workstations running in the cloud. Pixitmedia provided user and file storage via a PixStor cluster. All workstations sent over-the-shoulder feeds to a Moxion workstation for collaborative review.

Pre-visualization was powered by Epic Games Unreal Engine 5 and editorial workstations featured Adobe Premiere Pro. Additionally, a preview monitor was connected via a Streambox integration, allowing an editor to have a professional monitor on their desk. For compositing and VFX, Autodesk Flame's newly released integration with AWS Cloud Digital Interface (CDI) was showcased. Finally, for finishing and quality control, AWS partnered with FilmLight to demonstrate the Baselight Color in the Cloud solution and Colorfront for master QC.

At Colorfront, our goal is to be at the forefront of innovation and technology in the Media and Entertainment industry. Participating in the AWS IBC 2022 end-to-end production demonstration using our Colorfront Streaming Server technology with CDI-implemented support is a testament to our commitment to delivering the highest quality tools to our customers. We believe that this involvement will help to set a new standard for the industry, and we are proud to be a part of it.”

BRANDON HEASLIP, DIRECTOR OF SOLUTIONS ENGINEERING, COLORFRONT

BENEFITS

This implementation created a ‘single source of truth’ for media files in the cloud which all other processes can work from. As content leaks can occur from physical movement of media and hand-off points during a production, this one change alone can help enhance security for content owners. With content securely stored in the cloud, applications coming to the content, and every member of a production centrally authenticated, the risk of a content security breach is minimized. By centrally storing assets and applying the principle of least privilege to them, data silos are reduced, which also provides financial efficiencies through de-duplication.

By automating the publishing of assets, sidecar files, and project files, productions can create process and operational efficiencies. This potentially enables post-production cost savings and enables productions to spend more budget with creatives and put “more money on the screen” to make better content in less time. Real-time iteration adds to these technical efficiencies by allowing creatives to focus on being creative and eliminating travel time, analog processes, and lengthy synchronous approvals.

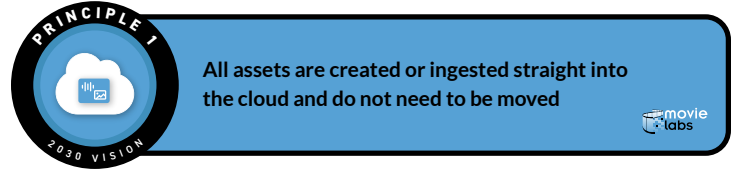
- The end-to-end workflow demonstrated operational efficiency by leveraging previsualization to help make informed decisions on-set, avoiding the need to “fix it in post”.
- In this workflow, uploading OCFs and proxies directly to AWS allowed editing and VFX to happen sooner than if the files were manually uploaded from hard drives at the end of the shoot.
- Leveraging cloud resources for storage and compute improved system reliability. For example, if a workstation failed in this instance, a new one would be spun up in its place within a few minutes.

ALIGNMENT WITH MOVIELABS 2030 VISION PRINCIPLES

The cloud architecture built by AWS and AWS Partners demonstrates a number of principles of the MovieLabs 2030 Vision that are technically achievable today.


PRINCIPLE 1

By ingesting media directly to the cloud, the production minimized data movement and confusion over which copy was the latest, where a specific version of the content was stored, and who controlled access to it. The single cloud core was connected to all downstream operations, applications, and processes and able to share assets and metadata with a central layer for access and management.




PRINCIPLE 2

Most Digital Content Creation (DCC) applications today cannot read directly from object storage, which results in assets being temporarily cached on cloud file storage when they are needed by the applications. For the AWS Studio in the Cloud workflow, these applications still came to the media by running on cloud workstations and the assets were hydrated to the PixStor cluster, providing the workstations with a file interface to interact with the assets. In the future, AWS envisions a world where DCC applications can natively read from and write to cloud object storage, removing the need for file system hydration..



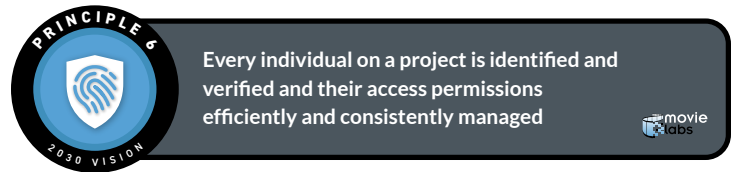
PRINCIPLE 3

All media resided in the cloud with Amazon S3 as the single source of truth for all assets. As assets were ingested into the cloud, event notifications were triggered to make the assets available to downstream consumers like AWS Lambda. The ability to make an asset available, i.e., publish an asset, is enabled by built-in notifications that are connected to the source of truth storage.



PRINCIPLE 6

Even though this was a closed production workflow, AWS served as the integration point across 16 vendors and demonstrated that a workflow can be secured with centrally managed authentication and authorization using AWS Directory Service and AWS Identity and Access Management (IAM) to uniquely identify participants and permission tasks and assets.



PRINCIPLE 10

This production utilized real-time compositing on-set of visual effects for both preview and director approval as well as final pixel rendering of the scene. These In-Camera VFX (ICVFX) demonstrate Principle 10 and how Virtual Production techniques can move offline VFX rendering (which can take days or weeks from asset build, rotoscoping and compositing) on to the set so the entire creative team can see, review and iterate on VFX elements in real-time on the set.



NEXT STEPS

While this demonstration of a holistic “studio in the cloud” provides a single architecture for an end-to-end remote workflow honoring the MovieLabs 2030 Vision Principles, more partners must be brought into the mix to truly provide content creators with the options necessary to enable the nuances of their pipelines. Such an environment requires agnostic interoperability between all components. AWS is committed to working with AWS Partners to increase the choice of applications and platforms that can be integrated and to eliminate compromises productions need to make to migrate to a similar environment. Future iterations will include more robust integrations with downstream media supply chain activities to extend the efficiencies created for delivery and distribution hand-offs.

With productions moving toward the cloud at an increasing pace, there is an opportunity to leverage cloud object storage (e.g., Amazon S3). Within media workflows, however, there are limitations because many applications expect a POSIX compliant file system – a system with a tree of directories, folders and files with names for each. This is very different from object storage systems where all files are stored in buckets. In the age of the cloud, this means either offering a mechanism to copy media to a compliant file system within the cloud environment, or presenting the object store through a translation mechanism, which often has performance and scalability implications. If common media applications were able to read and write directly from cloud object storage, productions could natively leverage the varied tiers and performance, durability, and security benefits of cloud storage.

MOVIELABS PERSPECTIVE

Delivering the 2030 Vision requires interoperable software platforms and change management, but neither of those can occur without a good foundation of secure cloud infrastructure. This AWS case study demonstrates that some of the toughest technical processes in modern media workflows – capture, color, compositing – can all now happen in the cloud without egressing data to on-premises locations for processing. Although this demonstration took place in a ‘controlled’ environment (in as much as any production is controlled), it demonstrates that the technical barriers to cloud-based production are falling. The next steps will be for partners to implement these solutions with real vendors, security systems, content workflows, and economic constraints on actual productions. As these systems improve and creative tools become more cloud-aware, we’ll be looking for less data movement (which is still happening in this example) and intelligent cloud systems like automated caching for performance. It’s encouraging to see how far we’ve come in such a short time, which bodes well for delivering new cloud native software-defined workflows.