

Moving to the cloud

A practical migration guide for Video Production, VFX and Post-Production Operations

Brought to you by



With sponsorship from:

RStor

February 2020

Introduction

In this paper we'll examine which workflows and processes are ready for cloud migration in 2020, how to cost effectively setup those cloud systems, and how to avoid the common pitfalls to ensure that your migration gives you the cloud advantages, without the cloud pain. We'll talk about the major Cloud Service Providers (CSP), their product set, business models and how to avoid mistakes that early adopters have made.

It is worth acknowledging that the promise of the cloud – freedom from local IT management, dynamic scaling of resources to match loads and OpEx replacing CapEx – are all very real and being enjoyed across multiple different industry verticals. Indeed, the *delivery* of media content to consumers leads the charge in cloud services with Over The Top (OTT) streaming services rapidly gaining subscribers that are migrating from traditional Pay-TV. Our focus here however is on the *creation* of Media and Entertainment content and that has been much slower to migrate to the cloud due to file sizes, incumbent investments in on-site hardware and the distributed nature of the workflows and decision making. There are considerable benefits in the cloud workflows, but there have also been many early adopters who have got their fingers burned with costs were that way out of budget, technology that did not deliver and unpredictable impacts on workflow.

As the costs of cloud storage and processing have declined, the availability of virtualized services has increased, and major production companies, including the Hollywood Studios, can now see that 2020 will be a key year in the cloud migration. For new Post/VFX providers that are born without the depreciating fixed costs of legacy equipment, the cloud has been warmly embraced. We are now seeing vendors that are 100% cloud native, that can scale their resources to customer demands and startup with little or no upfront hardware investment. These new startups represent a significant challenge to established incumbents that will need to match the innovation speed of smaller and more nimble “cloud boutiques”.

The Cloud for M&E Creation in 2020

In this section we'll evaluate the current state of the art, what is working in cloud today and what factors you should consider before you start. These use cases are assembled roughly in order of ease of deployment, value creation in cost savings or workflow efficiencies.

1. Archiving

Preservation of largely digital assets via magnetic data medium, stored in vaults -- co-located, or often within a few miles of the master files -- is a pretty dangerous redundancy option. The cloud offers an immediate solution to store massive volumes of data at very low cost, in locations far away from the earthquakes, hurricanes, national security threats, and other dangers of our major media creation markets. In today's modern and fast-moving media markets, digital archives should be available on the cloud for monetization, reference, and globalization, period.

There is almost no reason today why media archives should be stored on-premise when an always available cloud media library is likely cheaper and is certainly more useful.

2. Camera Original Assets

In many productions, proxies and camera original files (e.g. RAW) are created simultaneously but follow different paths to get back to the post facility, studio, and production offices. In most sound stage facilities, or even on location-based shoots, there are internet connections available for transfer of proxy files. The master camera files, however, tend to be aggregated on to hard drives and shipped back to post. This route introduces shipping uncertainty, security concerns, time delays and increase the chance of human error in re-associating the two types of files back together again. By providing extra bandwidth at shooting locations and a change of workflow to enable trickling of all files up to the cloud, files can be continuously synchronized to the cloud during the shooting day. Available bandwidth can also be saturated 24/7. A small piece of software may be required to prioritize the traffic based on the directors preferred shots but eventually, all files can be uploaded and be available for VFX vendors, editorial or marketing teams.

3. VFX Rendering

Bursting of large render jobs onto cloud Virtual Machines for overflow beyond an on-premise render farm is already commonplace with companies like Conductor (www.conductortech.com) specializing in managing the process. Some major movies and VFX houses have now moved to 100% cloud-based rendering, and yet, cloud rendering is not a simple job for your on-call IT helpdesk to configure. The complexities of mounting hundreds, or even thousands, of virtual machines each working on a fraction of a render job and managing those machines through jobs that can take a full 24 hours to complete is a specialist software task. Recompiling the returned output of the render requires a good traffic manager or outsourced render service to manage the process. However, the ability to say 'yes' to multiple jobs, knowing that the work will no longer be constrained by on-premise hardware, is an attractive proposition for VFX vendors.

4. WiP and Global Collaborations

The post and VFX world are becoming increasingly global as productions chase tax credits, as well as low-cost and available creative talent. Many of these locations become satellite offices but still require build-out of a full IT infrastructure. Multiple smaller offices add network complexity with private point to point links and a need for local IT support staff at each. In VFX, these problems are rife with up to 10 vendor locations collaborating on some major AAA movies. Each location stores local copies of the same files, runs file transfers between locations and other vendors on the project, and back to the client location. Each step incurs time delays, costs, and possible file loss or corruption. The cloud offers the ideal solution by storing all the files once (twice or thrice for resilience) and allowing each location permissioned access to that 'single source of truth'. Large files may need to be cached closer to the end user, but importantly they do not need to be replicated by each office location – the caching can be handled automatically in software by the cloud.

File transfer delays, rogue copies, and lost files are eliminated immediately, and if provisioned correctly, security is improved with additional features such as audit trails and version control enabled.

5. **Playout and Origination**

While not of direct interest to production and postproduction companies, the fact that the cloud is now an optimum playout system is relevant as shows made in the cloud are easier to distribute in the cloud. Many regional, national, and global distributors are now looking to the cloud to aggregate their programming so it can be easily versioned, packaged, and streamed as linear channels in multiple markets simultaneously. For post-production companies this is an opportunity to not only create the media, but also provide value added services in preparation and distribution of those delivery file packages.

6. **Regionalization Services**

This is another global use case as many vendors are often used - some in regions where owners would not want copies of their content available on local storage - and the file sizes are relatively small. A watermarked stream of the last edit of the content can be made available to captioning and dubbing vendors for time limited period. Many of these services, or at least an early pass of them, are moving to AI based translation. Those services are best run as serverless, cloud based services, where having the master files in the cloud makes the process much more seamless and avoids version control issues where regionalization vendors are operating on old edits because someone forgot to notify them of the change.

7. **Final Post Services**

Some of these are the hardest, and most likely the last to migrate, as they often involve custom hardware or absolute color critical displays with very low latency. We may not yet be at the point where a remote cloud machine can stream a color-critical 4K master that can be scrubbed and power-windowed in final color. However, the final composite render could be handled in the cloud once the CDL and EDL are completed and uploaded.

M&E workflows are complex and there are opportunities for most use cases – at times with some modifications -- to be migrated to the cloud in 2020. Now, it is worth looking at the common pitfalls and problems that can occur in a cloud deployment.

Considerations with Cloud Deployments today

As with all new technologies, there can be some initial deployment issues and problems. Fortunately, enough companies have now explored these use cases that we can all benefit from their experience. Here, is a list of the major issues that M&E cloud deployments have encountered, and in Section 3 we will explore how they can be mitigated by design.

The most common issues with the cloud today are not the technology, but the business model which is sufficiently different from the current capex and depreciation models, to cause significant issues if not managed effectively.

Here is a breakdown of problems and the cause of each:

Bidding against yourself	Much like Uber's "surge pricing", organizations may pay more for cloud compute resource if working in a popular market, at a particular time of day, and/or their cloud provider is running low of resource in that market. Capitalism, at its core means the variable 'spot' prices will increase. If you don't cap your expenditures or set alerts, you could see prices increase beyond your budget. This can be particularly acute when doing massive render jobs where even a fraction of a cent in machine cost can add up very fast and cost overruns can happen in hours.
"Just Hit Render"	VFX artists are used to a render farm on premises, which, albeit limited in capacity, is largely considered to be 'free' to use. It isn't free, but those artists don't see any direct cost, so it's perceived to be free. That approach can become very dangerous when there are direct costs for every render run and those costs are not visible to the artists.
"App not available on this cloud"	We're seeing a fracturing of major applications across the different cloud service providers such that there is no guarantee today that your app

	<p>will be available on the cloud you select. When coming from the relatively stable world of PC media tools where most everything is available for Linux, Mac, or Windows, this feels like a step backwards.</p>
Tiered Storage	<p>Although billed as an advantage with cloud services, there can be some nasty surprises with tiered storage. These issues vary widely between service providers and the costs vary by the accessibility of your content. For 'colder' storage you'll pay less per GB stored but significantly more when you want to access your media, with charges and time delays to 'warm' the files back up, then more charges to move it to a hot enough tier that you can actually work on them.</p> <p>To truly take advantage of the different tiers you will need to thoroughly understand your storage access history over several months in order to generate a good usage model before selecting tiers. Being forewarned is being forearmed.</p>
Pricing varies by region	<p>If you've run your own datacenter, you'll know that space, power, and cooling vary between cities and countries. Major CSPs have, have the same issue and pass those costs on to customers. This requires you to understand the specific regional costs in your own model. Pricing for the same services can vary by region by up to 100% on storage and 50% on compute. Also note not all services/machine types are available in all regions.</p>
Access and Egress fees	<p>Your hard drive doesn't charge you to access your data, but the major CSPs charge for data egress sometimes and access fee to 'wake' the content up from deep storage. Again, you need to understand your data usage to ensure this will</p>

	work for you as the 'tax' on accessing your media can be large and hard to predict.
API Charges	<p>Your hard drive doesn't charge you to send a read or write command to it, so once again this one can be a surprise. Although small (fractions of a cent), these costs can start to mount when writing thousands of small files in and out of the cloud. Each API request can be sent by your MAM or application as a separate data call for which you will be charged, although you may not directly see them.</p> <p>This is a very hard charge to model in advance as it's difficult to know how often your applications make a data call.</p>

Table 1: List of common economic issues with major Cloud Service Providers

To understand how these costs add up, here is an example of an archive use case where media stored in deep archival storage in one region (for safety it is away from the primary, assume central region) is pulled to a West Coast region for processing for 1 month (so exists then exists in two cloud regions), and then the original assets are returned to the archive region. We'll consider two options for the how often these assets are pulled – one assuming pulling 5% of the archive and one assuming 20% of the archive is pulled back. As you can see, the costs can vary considerably (by over \$15,000) if you underestimate how often you will reclaim the media from the archive.

Cloud Archiving Scenario uses typical cloud costs (accurate as of Feb 2020) for the month in question.	Assume 5% of archive accessed ¹	Assume 20% of archive accessed ²
Cost of cold tier archive storage (\$0.0045 per GB)	\$4,719	\$4,719
Cost of API to read data (\$0.011 per 1000 requests) from cold storage.	\$0.57	\$2.31
Cost to move media to 'hot' storage tier (\$0.033 per GB)	\$1,730	\$6,921
Cost to move media to Western Region on CSP's backbone (\$0.02 per GB)	\$1,049	\$4,194
Cost of media storage in hot tier for 1 month (\$0.024 per GB)	\$1,258	\$5,033
Cost of API to read media from hot storage for compute processing (\$0.0044 per 1000 requests)	\$0.02	\$0.09
Cost of API to write media back to hot storage (\$0.0055 per 1000 requests)	\$0.29	\$1.16
Cost of Intra-region transfer to move processed media back to central region (\$0.02 per GB)	\$1,049	\$4,194
Cost of API to write media to cold storage (\$0.055 per 1000 requests)	\$2.86	\$11.55
Total Costs	\$9,808	\$25,076
Costs related to actual storage	\$5,977	\$9,752
Costs of other charges	\$3,831	\$15,324
Other charges as a % of total	39%	61%

Table 2: Illustration of complexity of cloud costs from major CSPs for a typical cloud media scenario during one month. Considers just the cost of cloud storage, not any cloud based compute services on top. Also demonstrates how non-storage related charges can account for over 60% of monthly costs

¹ Equates to 51.2 terabytes and 52,000 file accesses, if each file is 1GB

² Equates to 204.8 terabytes and 210,000 file accesses, if each file is 1GB

Lessons Applied

After all the doom and gloom, it's probably worth mentioning again that the cloud is a great and effective way to deploy services if you go into it with your eyes wide open. Let's look at some best practices to ensure you don't fall into these common pitfalls.

1. Over-Specify the Network

Providing bandwidth is like building lanes on the 405 in LA, or the M25 in London – they'll be filled as soon as they are built! One of the biggest inefficiencies our industry has is the time wasted while waiting for transfers to occur, or for files to be uploaded/downloaded. As the old adage goes, time is money. Saving a few cents on bandwidth can cost multiple dollars on wasted time later. Therefore, it's worth over-specifying your network connectivity – both the local loop to your regional POP, and the onramp to the public clouds. This includes not just facilities, but also production studios themselves who used to get by ordering a T1 connection from a remote location. That will no longer cut it in the modern cloud world.

If you think you'll need 1GB, order 10GB circuits. If you think you'll need 10GB, order 100GB. Trust us, you'll use it.

2. Watch for network overlap and consider network separation

There's nothing like provisioning 50 VFX workstations remote streaming from the cloud, then realizing they slow to a crawl and become unusable when someone sends you a big file. Streaming desktops, VoIP and accelerated file transports rely on discrete IP protocols and should be separated on VLANs. If you do not manage your network effectively by separating the traffic, and specifying Quality of Service, you will find your latency and bandwidth sensitive use cases failing.

3. Implement user behavior change control

As noted earlier, it's easy to incur cloud costs because your users don't understand the costs associated with its usage. Beyond user training, you may want to expose cloud costs down to the supervisor, or even end-user level so the team can understand the ramifications of their cloud activities. One of the

benefits of the cloud is the ability to instantly turn on resources dynamically but users need to understand that meter is always running. It may also be appropriate to put a "cloud dial" app on their desktop with options for 'get this work done overnight for cheap' or 'this is urgent work and I need it right now'. The prices could vary wildly between the two options, and when the director is leaning over your shoulder and approving the spend, it's worth turning the dial up to 11 to get the work done right now. In many cases that is not required, and devolving the power down to the users will help them to understand the variance and be empowered during the transition.

4. Rethink your workflows

This is a biggy. Read the MovieLabs 2030 Vision Paper (available for download free from www.movielabs.com/production-technology) then think about the user and applications coming to the data. Today we tend to send the media to the users and to the application that they are using to do the work. That is fairly non-sensical when the applications and their license files are tiny by comparison to the size of the media files. MovieLabs predicts a future where the workflow is turned on its head and the users can remotely access media in the cloud via virtual workstations.

Many of the services in post-production occur bringing media in and out of the physical facility - that enables the facility to charge extra fees (for ingesting files, quality checking them, pre-processing them and then reversing the process for delivery back to the client). What will happen to those services when the content owner may just provision the post facility access from a central cloud store? This may mean less people doing basic 'data wrangler' jobs, and more focus on creative employees doing their work, remotely on the assets stored in the cloud. This could have radical implications for the way service providers hire and develop staff, or even how large a physical office footprint needs to be. Think about what services you could offer as purely virtual cloud-based micro-services. Could you have your services available globally via a post facility cloud marketplace? Could you develop AI services based on your specialist skills and train those models with your own staff?

The implications and change the cloud will enable are considerable. All companies need to start preparing now, as it's certain your competitors are.

5. Rethink how applications and plugins are built

Most applications used in Media and Entertainment were built to run on local machines or servers using a file system and file storage which is local to them. Those applications have been extensively modified over the years to be adapted to each organization's unique workflows. While those applications can be 'lifted and shifted' to run on a virtual machine in the cloud and therefore be "cloud capable", that is not the same as a truly "cloud native" service. By simply running PC or Linux applications in the cloud we're missing the opportunity to deliver on many of the cloud benefits.

Cloud native applications are built to take advantage of the scaling nature of storage, compute and parallelizing those functions – while appreciating the inherent cost complications of the cloud. Cloud native applications can seamlessly scale to multiple concurrent users, have licensing models which scale linearly and can be extended and combined through extensive micro-services and APIs.

While the major software vendors work to re-architect their products to be truly cloud-native, each organization needs to reexamine the customizations they have applied to those applications and assess whether they are still relevant or should also be rebuilt. That may also combine elements of points 3 and 4 above and use this as an opportunity to separate users from the 'comfort blanket' of file names and file trees for asset management. A modern movie can have millions of individual media elements and for them to be arranged in a file structure and described via intricate combinations of codes embedded in their filenames is a system which has met its end. By rearchitecting around modern metadata schemas that are linked to the files they describe, we can enable much richer asset management, discovery and workflow orchestration systems. We can also avoid the frequent problem of "we can't lose Bob – he's the only who knows where the files are kept"!

6. Build your own cost calculator and model your data

As explained in section 2, the most common cloud complaint is 'sticker shock' from not understanding how the costs line up with your particular data models. There are a few calculators available that cover the costs of the major cloud providers but it's worth noting that the CSPs change their pricing multiple times per year and frequently just within certain global regions. Keeping these calculators up to date is a major challenge. It is best to assess your own use

cases, data flows and compute requirements, then build your own multi-cloud price comparison system. This system should consider the storage, compute, and networking charges from each cloud provider.

One solution is to pick one cloud provider. This can simplify the calculator issue considerably, but this leads us to...

7. Don't just pick one cloud provider

While it's tempting to make life easier to just pick one provider and stick with it, there are many reasons why it is a bad idea in the medium and long term. For a start, the major cloud providers are still in a period of rapid growth -- adding services and regions all the time. You may pick one today because you like their combination of compute options and AI services, then find next month one of their competitors trumps them with new or improved products.

By picking one provider you also limit your negotiating power because they know you are unable to easily move your media, as they charge you to do so. Also, by locking your media to one cloud you are beholden to them, which limits your ability to serve clients across all the major clouds.

We are expecting a true multi-cloud future where every media application is available on every cloud compute service, but unfortunately, we are not yet there. Therefore, you need to be prepared to use the compute services from different clouds to build a workflow that includes all of the applications/services you use today, and potential new serverless cloud functions which may spring up in the future. There is another way to stay nimble and that is not to put your media in any of the major CSPs and ...

8. Consider RStor

RStor (www.rstor.io) provides cloud storage services that are data independent from the big cloud providers, with a model that avoids every single add-on charge in Table 1. RStor doesn't charge any fees for putting your data into the cloud (ingress), accessing or processing it (API and data access charges) or taking it out again (egress). RStor also charges the same price in every data center worldwide and offers one tier of service (effectively the speed and access of S3 but at Glacier pricing) which makes price calculators much easier.

RStor built a network with super-fast connections to each of major cloud service providers, which means you can keep your files in RStor and use CSP's for compute services. You can store 3 copies of your files (say one in London, one in East Coast, and one in West Coast) and using RStor's network backbone have automatic backup between the 3 copies which are always in sync, for no additional charge. This way you avoid data transfer charges from the CSPs when you upload files in London, because the files are automatically synchronized to Los Angeles, and vice versa.

The net effect is that you can store all your media in RStor, and then broker compute across all of the CSP's. By keeping control of your media, you stay nimble to future innovations without limiting yourself to just one. RStor is available through media and entertainment specific service providers such as Sohonet's Filestore+ service as well as direct, all at the same low price.

SECTION 4. Conclusions

The Media and Entertainment post and VFX industries have grown considerably in the last 5 years due to the explosion of new TV content, and the increase in complexity of VFX. However, we've come to an inflection point where we simply can't add more capacity in many cities (e.g. the talent is simply not available or has become too expensive.) Our only real option is to do more with what we have, to drive innovation in workflows and efficiencies, and create more content with the same or less resources. The cloud has matured at just the right time for us to take advantage of the technology to initiate those changes.

We're entering an exciting world filled with innovations and the ripple effects of the cloud could mean changes not just in technology, but also in the economics of our industry. This disruption is beyond anything we've seen when we moved from film to tape, or from analog to digital. The potential of the cloud could upend not just the job functions in our industry, but where those jobs are done, and what skills are required to do them. The winners will be the organizations that can embrace that change and be willing to redefine tools and workflows to be cloud native as purely 'lifting and shifting' the existing processes will not leverage the advantages inherent in the cloud.

While it's highly unlikely AI or robots will ever fully replace the creative roles within our industry, cloud functions today can automatically operate on files without human intervention and relieve many of the mundane and repetitive tasks (and associated fees) of file processing. We could certainly see an end to manual tasks such as ingest, conform, transcoding, QCing files, and even rotoscoping, as automated services could start taking over these functions. Expect inefficiencies to be eliminated from workflows where teams wait for files to be delivered before they can start working. In fact, it may be worth investing now in a new scheduling service that can intelligently understand the work every member of the team is doing, and prepare their next job in advance by prepositioning their files, loading their applications and preferences and scheduling their team - so that everything is ready for them to start their creative work.

Post and VFX companies have always adapted to new technologies, and it's clear the nimblest will win as we navigate to this new cloud world. We look forward to sharing the journey with you.



Headquartered in the Los Angeles Underground

Try to find us @ www.entertainmentconsultancy.com