Integrating cost and time metrics into media asset management workflows: Understanding the true cost of cloud storage for media archives

Received (in revised form): 13th January, 2021

Roy Burns
Vice President of Media Solutions, Integrated Media Technologies, Inc., USA

Roy Burns is Vice President of Media Solutions at Integrated Media Technologies. As a member of the advanced consulting team, he partners with clients to map their current and future media workflows, define best practices, and recommend and implement scalable forward-looking media asset management solutions. He is a named software patent holder for 'Data Migration Management and Migration Metric Prediction' and is on the forefront of emerging media workflows and technologies. He is a graduate from the Musicians Institute's Recording Institute of Technology and an active member of the Society of Motion Picture and Television Engineers (SMPTE) and the Hollywood Professional Association (HPA).

Integrated Media Technologies, Inc., 5200 Lankershim Blvd, Suite 700, North Hollywood, CA 91601, USA
Tel: +1 323 791 6481; E-mail: roy.burns@imtglobalinc.com

Abstract Metadata plays an important role in media asset management (MAM), powering access control, searchability and workflow automation. The more we know about our assets, the more efficient and usable our digital asset management (DAM) and MAM systems become. As cloud and operating-expense services have become more ubiquitous, it has become essential to understand time and cost implications and how they correlate to media workflow. This article discusses the value of tracking cloud storage and egress metrics as metadata within DAM and MAM systems. By correlating these metrics against existing metadata, administrators are able to build total cost of ownership models, budget more effectively and enable a 'chargeback' usage model. These metrics also help to manage end users' expectations more effectively within the MAM and storage ecosystem.

KEYWORDS: metadata, SLA, MAM, cloud storage, TCO, business intelligence, budgeting, chargeback

INTRODUCTION

This article references use cases within the media and entertainment landscape, the borders of which are increasingly nebulous. It presents viable knowledge and reference architecture that will benefit any company large or small, working with media. While larger companies will have more resources than smaller ones, at the end of the day, everyone cares about the bottom line: the more light an organisation can shine on resource justification, the better it is for everyone involved.

The media workflow pipeline begins well before the cameras roll and ends long after the finished product is delivered. Some companies own that pipeline from end to end, while others are passed the baton midway through the process. Recognising that the range of tools to empower this pipeline is extensive, this article focuses primarily on the storage and asset management components.

Not everyone needs to edit four streams of 8K RAW, but higher-resolution and larger bit-rate file types are increasingly common. Sizing editorial storage is a key component to any media workflow. Understanding the basics of capacity and bandwidth are the primary factors. Creative users do not want to think about the storage and networking back end — that is not their job. Rather, it is the storage administrator's job to facilitate the creative end user's work, minimising friction on their behalf. End users need the storage to be performant, and they need to be able to search effectively and find the assets they need in a timely manner. It is also the administrator's job to control costs and budget accordingly. After all, it would not be cost-effective to have petabyte-scale all-flash storage. This is accomplished by tiering data to lower-cost/lower-performance storage volumes.

The criteria for placing data on the various tiers of storage vary, as does the method of data movement. Whose job is it to move data from the various storage tiers, and what is the best tool for that job? Creative users are not data wranglers, and were it up to them, they might just ask for that petabyte-scale NVMe array. Storage administrators can move data manually or employ automated methods to tier data based on various criteria, such as access time or data modified. This, however, risks upsetting the end user if they are unable to find what they need because it is offline or it is taking too long to bring back online. How can companies address this point of friction?

Media asset management (MAM) systems are the bridge between the administrative functions and the creative end users. With rich metadata, custom workflows and automation, MAM systems empower end users to perform self-service wherever possible. A well deployed MAM system will have proper roles-based access control, workflow automation and metadata galore.

The topic of metadata is broad, and there is no 'one-size-fits-all' taxonomy that applies to every company. What is relevant to some is not relevant to others, even within the same organisation. At a base level, a MAM system should capture technical metadata about the assets housed in its system — things like codec, container, frame-rate, bit-rate, etc. Beyond technical metadata, descriptive metadata can be used to describe the essence of the data — things like Project/Job Code, or Show/Movie/Sports Event. More recently, artificial intelligence/machine learning (AI/ML) have come into favour, enabling automated metadata extraction of things like automated speech-to-text, object detection, facial recognition, and much more. Such AI/ML generated metadata will be time-based, and often reflected as markers within a timeline in the MAM system as well as the non-linear editing (NLE) application.

With this enhanced metadata, users can search more easily for what they need and pull it into their project accordingly, but what about archived media? MAM systems often integrate with a third-party data mover to empower self-service and restore jobs to end users. At present, the biggest hurdle is managing the end user's expectations. The MAM system will keep a low-resolution 'proxy' version of all assets online all the time to provide users with a point of reference when searching. These lightweight proxies are also often used in a proxy editorial workflow. At some point though, the project will need the high-resolution source files back online. The means by which this is done is what needs to evolve.

Administrators need to understand the cost implications of storing assets across various storage tiers, and users need to know how long it will take before they can work with assets — neither of which are unreasonable requests. Where the archive sits is often unclear to the end user, who does not care, even though where those archived assets sit very much affects them. There are numerous tiering strategies that companies can employ, including high-density tier-2 disk arrays, LTO tape, private object storage and public cloud storage. This article focuses on cloud-based archiving and the unique complexities that are introduced with the operating-expense nature of public cloud storage.

CLOUD STORAGE

The shared responsibility model of the cloud has shifted many of the administrative overhead costs associated with storage away from capital expenditure in favour of on-demand 'pay-as-you-go' operating expenditure. The elasticity, durability and availability have made cloud storage an attractive archive target for media workflows.

Cloud service providers (CSPs) bring a lot to the table beyond storage. It is increasingly common to see compute and transformation workflows take place entirely in the cloud. MovieLabs — a nonprofit research lab run jointly by Paramount Pictures, Sony Pictures Entertainment, Universal Studios, Walt Disney Pictures and Television, and Warner Bros. Entertainment — predicts that in ten years, all assets will be created or ingested straight into the cloud and will not need to be moved.[1] Based on the current trajectory and rapid adoption, this is certainly the way things are moving. More recently, the Hollywood Post Alliance proved this out by creating a short film entirely in the cloud, validating the future concept as something that can accomplished today.[2]

From a storage standpoint, there are various classes broken into two main types. Block storage is the most performant class, and by far the costliest. Block storage in the cloud is required for an editorial workflow housed entirely in the cloud, ie a Teradici-style PCoIP cloud edit bay. Block storage provides more flexibility as object storage does not let the user change parts within the object. Block storage options are definitely not an archive tier, but they can be a restore target, which is why they are called out here.

For media archives, companies leverage object storage. Amazon AWS S3, Microsoft Azure Blob and Google GCS are the main players in the public cloud space. Each CSP has various subclasses within its storage that determine availability and durability. For example, AWS has six storage classes within its S3 family, the most inexpensive tier being 'Glacier Deep Archive' with a list price of US$0.00099 per GB/month, or about US$1 per TB/month, and the most expensive tier 'S3 Standard' ranging between US$0.021 per GB/month to US$0.023 per GB/month or around US$23 per TB/month.[3] With such a massive cost variance, the cost of a company's archive may swing drastically. Optimising what storage belongs on which tier is a delicate balance because each option has cost and time implications.

As Table 1 shows, there are a number of variables. The first is durability: each class features '11 9's' of durability, meaning that if a company used this solution to store 10,000 objects, on average Amazon might lose one of them every 10 million years.[4] Next to consider is availability. Within each geographic region, Amazon will have multiple isolated data centres referred to as 'availability zones'.[5] Broadly speaking, these function like redundant arrays of independent disks for a geographic area. Every tier, with the exception of S3 One Zone-IA, will geo-spread data across at least three separate physical locations or 'availability zones' within a single region; this provides fault tolerance within a region. Figure 1 illustrates how data may be stored in an AWS region.

S3 One Zone-IA is not a recommended tier for archival storage. Best practices suggest skewing either to S3 Standard-IA or to S3 Glacier, depending on what best suits the company needs.

The cloud introduces new metrics that companies have never had to consider for on-premise/capital expenditure storage; these include:

Table 1: Performance across the S3 storage classes

	Standard	Intelligent tiering	Standard-IA	One Zone-IA	Glacier	Glacier Deep Archive
Designed for availability	99.99%	99.99%	99.99%	99.95%	99.99%	99.99%
Availability service level	99.99%	99%	99%	99%	99.99%	99.99%
Availability zones	≥3	≥3	≥3	1	≥3	≥3
Minimum capacity charge per object	N/A	N/A	128 KB	128 KB	40 KB	40 KB
Minimum storage duration charge	N/A	30 days	30 days	30 days	90 days	180 days
Retrieval fee	N/A	N/A	Per GB retrieved	Per GB retrieved	Per GB retrieved	Per GB retrieved
First byte latency	Milliseconds	Milliseconds	Milliseconds	Milliseconds	Select minutes or hours	Select hours

Source: Amazon (2021) 'Performance across the S3 storage classes', available at: https://aws.amazon.com/s3/storage-classes/

- *Minimum storage duration charge:* With all this talk of cost and storage optimisation, it is not uncommon to conjure ideas of moving data between tiers on a daily basis. This would actually be the *opposite* of cost optimisation as it would drive cloud spend way up. To keep prices low for the non-standard S3 tiers, CSPs build in some guard-rails. If a company stores content on infrequent access or glacier tiers, it will see significant cost savings. One of the ways the CSPs keep costs low is by building in protection — they do not want administrators storing content in these tiers for a few days only, so it is important to be aware of the minimum storage duration when leveraging one of these cost-effective tiers. If content is moved or deleted from one of these buckets before the minimum time, a penalty will be incurred and the organisation will still end up paying the same as if it were there the entire time. For example, if a storage administrator put a 52 GB file inside of S3 Deep Glacier for 60 days and then moved it out, they would still be liable for the remaining 120 days, as specified in the minimum storage duration.
- *Retrieval fee:* This requires some additional context as there are multiple variables to consider. If someone moves content from a glacier bucket to S3 standard, they will be charged per GB. If someone needs to retrieve that asset back to on-premise storage to a server message block (SMB) or network file system (NFS) target, they will also be hit with egress fees on a per GB basis. These fees, combined with minimum storage duration, have proven to be complex hurdles for customers who have a difficult time quantifying the amounts of data they plan on moving in and out in a given year. For data coming out of S3 Glacier, there are even more retrieval options:
  - *Bulk:* As the name implies, this is the cheapest method for retrieving large transfers. It comes at the cost of time, however, with jobs at this service level taking an average of 5–12 hours (based on 'time to first byte' — see discussion of first-byte latency, below).
  - *Standard:* Mid-range or typical retrieval will average 3–5 hours, but will likely cost around 12 per cent more than bulk retrieval.

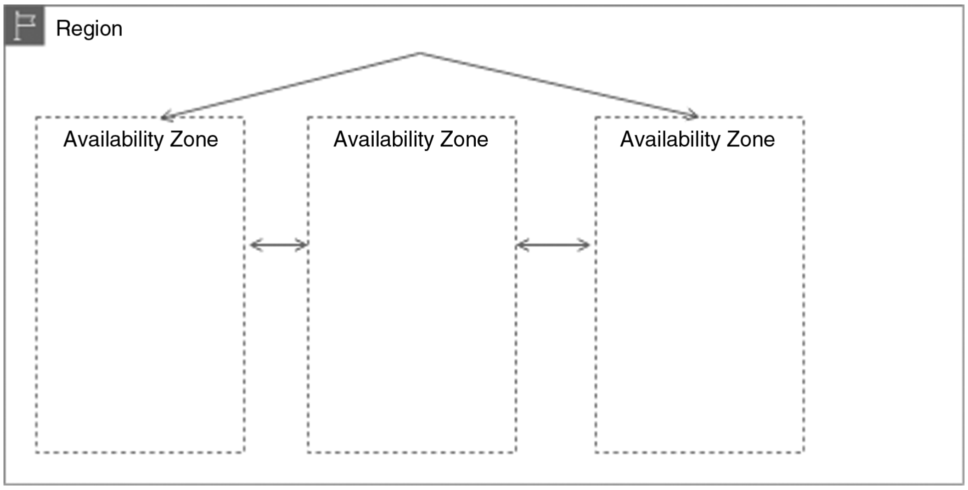

Figure 1: Fault tolerant availability zones in a single AWS region

- *Expedited*: The most expeditious option is also the costliest. With expedited retrieval, it is possible to get one's data within minutes, but one should expect to pay roughly 21 per cent than with bulk retrieval. AWS Glacier Deep Archive is an even cheaper tier, but does not allow for different retrieval service levels. Estimated retrieval time from this class ranges between 12–48 hours. So, while it may be the most cost-effective, it is essential to manage expectations on how quickly users can expect content to be returned to them.
- *First-byte latency:* This variable does not pertain to cost estimation; rather, it speaks to how long it will take before Amazon will be able to read the *first* byte of a given asset. The first byte of a given file is simply the lowest common denominator. By way of example, there is a large difference between a 2 MB spreadsheet and a 187 GB MXF Master: in both scenarios it would take the same amount of time to reach the first byte — or start — of the asset, but due to the size variance, the larger file would take longer to restore. While first-byte latency means little to the storage administrator, it means even less to creative end users, who want to know how long it will be until they can actually work with the asset(s) in question.

To get an idea of how this all looks, see Table 2. The time estimation here is larger than what is referenced in the bulk, standard, expedited descriptions above. This is because the size of the job affects the transfer time.

With these three retrieval service levels, cost and time estimations can easily balloon if left unchecked. An expedited retrieval is up to 300 per cent faster than a bulk retrieval, but costs 35 per cent more, excluding any 'minimum storage duration' penalties that may be incurred.

As complex as these variables may seem, with the right level of wrangling, a company can provide its end users with a cost-effective, highly available archive.

Azure Blob storage is subject to the same variables noted above, and the penalty for early deletion is the same (180 days). Despite having fewer classes, the same principles apply to its Hot, Cool and Archive blob tiers. Also, in place of three retrieval service levels, Azure has two: 'Standard' and 'High Priority', with high-priority retrieval being a 400 per cent cost increase over standard. It should be noted that Azure also provides a discount for reserved capacity. For predictable storage over a 1–3-year commitment, Azure will (at the time of writing) discount the cost up to 20 per cent at 1 PB scale.

Table 2: Job estimation

	Job size (GB)	**Estimated time** (hours)	**Estimated cost** (US$)	**Cost delta** from bulk (%)	**Time delta** from bulk (%)
Bulk	187	16	17	0	0
Standard	187	9	19	12	78
Expedited	187	4	23	35	300

Note: Costs used in the example are based on the AWS list price for us-west-2 at time of writing.

Large customers will often negotiate custom terms with CSPs that take into account various pricing concessions. For everyone else, the list prices provide as a framework to better understand cost modelling.

CLOUD-TAGGING

Cloud service providers provide a robust method of tagging compute and storage resources for a multitude of use cases, including:

- cost allocation;
- automation;
- access control; and
- resource group organisation.

Tags act as metadata that are assigned to cloud resources in the form of 'key-value pairs', defining a key and accompanying entry for that key. When it comes to assigning resources to cost centres, this key-value pair method is essential. Where it can come up short, however, is in tagging S3 buckets and S3 objects. The amount of S3 objects will easily dwarf the amount of compute resources — manually assigning a tag to every object manually would be nearly impossible.

Companies can glean automation by applying a key-value pair to an S3 bucket,[6] and an administrator can define a lambda function to copy the tags from the bucket to the object. This is a decent start, but it assumes the tags on the S3 bucket are enough to drive the relevant automation needed. If a company wants to automate data movement, control access, and understand who is responsible from a 'chargeback' standpoint, it would need multiple buckets with multiple policies. This is not impossible, but it is certainly very complex when paired with a MAM system that often references a bucket as a single repository for all assets.

It is also important to consider that tags on a S3 objects will act as metadata for the object itself, so what happens when transferred back to block storage, or when it goes back from block storage to object storage again? The complexities and administrative overhead of tagging objects compared against the relative low cost of storage has made many ignore the concept altogether, thinking the outcome does not justify the effort.

MAM RE-ENTERS THE EQUATION

As a front-end tool, the MAM system acts as the database of record for media assets. It is where all relevant metadata reside, with metadata associated at the asset level regardless of whether the asset is stored on block or object storage. As described earlier, the MAM system houses technical, descriptive and temporal metadata, all of which inform user access control, search and find, and workflow automation. Taking this one step further, one can correlate two new dynamic metadata values: cost and time.

These variables are dynamic because both cost and time are ever-changing. Rather than store them as a static key-value pair as described in the cloud-tagging section above, a workflow should run a multi-variable calculation and return the dynamic result to the user. How long will it take before a user can edit these assets? Is there a cost associated? If so, how much? Is it possible to bring back the content more quickly?

MAM systems are well known to be workflow orchestration engines — they communicate with various applications to automate the media pipeline, from third-party transcoders to AI/ML tools. One such integration is with data movers. Out of the box, a MAM system might natively support the ability to write to cloud storage targets for archive. For some companies, this is enough. Often times this is a barebones implementation relying solely on the standard S3 'put' and 'get' commands. This technically works — it will move data to and from the cloud — but it will fall short of the intended goal of providing dynamic insight to assets. The MAM system is a powerful tool — the building blocks are all there to help build conditional workflow logic. The following solutions can be used to accomplish more intelligent and contextual archive and restore jobs:

- *Native/built-in MAM support:* Using the native S3 'put' and 'get' commands, the MAM system can send and retrieve content from the cloud. Be aware that if a company relies on the cloud to auto-tier their data, that workflow will wait (or potentially even timeout) until the asset is ready. In this case, the workflow author should take proactive measures to understand where the assets currently reside and provide a high-level estimation of time and cost to the end user. Note that when relying on the CSP's 'auto-tiering' policies to move data dynamically, it is important to consider reporting this movement back to the MAM system. If the MAM system sends data to an S3 standard bucket, but the CSP subsequently moves the data to a cold tier, consider an automation policy to trigger a notification to update the metadata and status within the MAM system.
- *Third-party intelligent data mover:* Integrating with a dedicated data mover can provide benefits such as faster S3 multipart upload, multi-cloud support, writing directly to glacier or cool blob tiers, and the ability to select faster retrieval. These tools can also provide insight and automate movement based on technical metadata. By integrating into the MAM system, the MAM remains the database of record and the single pane of glass for users to interact with. Workflow authors can take a multi-variable approach of marrying metadata from the data mover, the MAM system and the S3 bucket to provide users with the dynamic output for archive and restore jobs. In some integrations, the ability to perform a 'dry run' can report to users how long and how much a job will cost before any data movement takes place.

The effort of integrating conditional logic into data mover MAM workflows ensures users have a better idea of how long it will take to retrieve assets. Managing expectations goes a long way even if it means delivering the news that a user will need to wait 12 hours for content to be available to edit. With proxy workflows, this often means an editor can start working on a project while the source media are restored in the background.

BUSINESS INTELLIGENCE

Providing a better end-user experience removes friction for creative users. Administrators also see value in the dynamic properties of 'time and cost' inside of MAM workflows. The addition of these two variables opens a new world of business intelligence outside of the MAM workflow itself. Cloud resource tagging is an effective method to correlate costs to users, groups and jobs. MAM systems throw a spanner into the works, as they are often a multi-tenant application. It is not uncommon for a MAM system to service the needs of multiple business units.

When the cloud bill comes at the end of the month, there is one line item for data storage and egress triggered by the MAM application. This is because there is a single API user that MAM system uses to interact with the cloud. Even with proper tagging, it all funnels back to this API user. The MAM system contains a treasure trove of context that can be correlated to users, groups and jobs. With the ability to correlate costs and time to existing metadata, companies can do internal 'chargeback'. IT departments have leveraged some form of chargeback and showback for years, with the rise in subscription-based services and cloud computing, chargeback mechanisms have become even more palatable.[7] In-house creative teams have also started to adopt this model; indeed, as an in-house creative industry report from Cella states, it is the natural evolution for an internal creative services organisation.[8] Agencies that need to associate costs to their clients have an even more critical need to understand cost of goods sold.

Understanding the value of cost and time correlation against existing metadata in the system can lead to some critical business intelligence. The following examples outline typical use cases with supporting sample reports:

- *'Chargeback' to department:* In this scenario, the MAM system is the portal for multiple departments to ingest new media and find legacy content to work with. Consumer marketing, affiliate marketing, production and social media departments all leverage the same central repository of content. In a given month, roughly 85 TB of content is ingested and archived to the cloud. In the same scenario, 15 TB are brought back from archive. The cloud bill for the 12th billing cycle shows just over 1 PB of capacity in total spread across S3 Standard (200 TB) and S3 Glacier buckets (800 TB), with 180 TB restored from Glacier over the course of the year. Without correlation, the cloud storage and egress bill would be split evenly among the four departments. With correlation, as shown in Figure 2, administrators can see that the production department is responsible for 80 per cent of storage utilisation as it ingests all of its raw source material. No other department has access to this raw source media. The finished outputs are also ingested into the MAM system; other departments *do* have access to these assets and use them frequently. The same correlation shows that the social media department is responsible for 90 per cent of all restores as it is constantly going back to the archives to post 'Throwback Thursday' retrospectives. Due to the cost implications of the various restore service levels, a report breaking down how often each restore type is triggered helps administrators understand how often the costlier 'expedited' retrieval is leveraged. Using time and cost correlation against technical metadata, it is possible to identify that the majority of assets retrieved from the archive were created in the past calendar year. Equipped with this knowledge, the administrator can make better informed decisions around sizing on-premise storage, as well as the necessary adjustments to the archive policy which may have been too aggressive.
- *'Chargeback' to job code:* Consider a scenario similar to the one described previously, but in this case, while the MAM system is still a multi-tenant portal, the departments work in concert with each other to fulfil projects for external clients. This is for a VFX + post finishing shop, so the departments look vastly different. This hypothetical company has motion graphics, craft editorial, colour and audio sweetening teams all using the MAM system. Without correlation, the

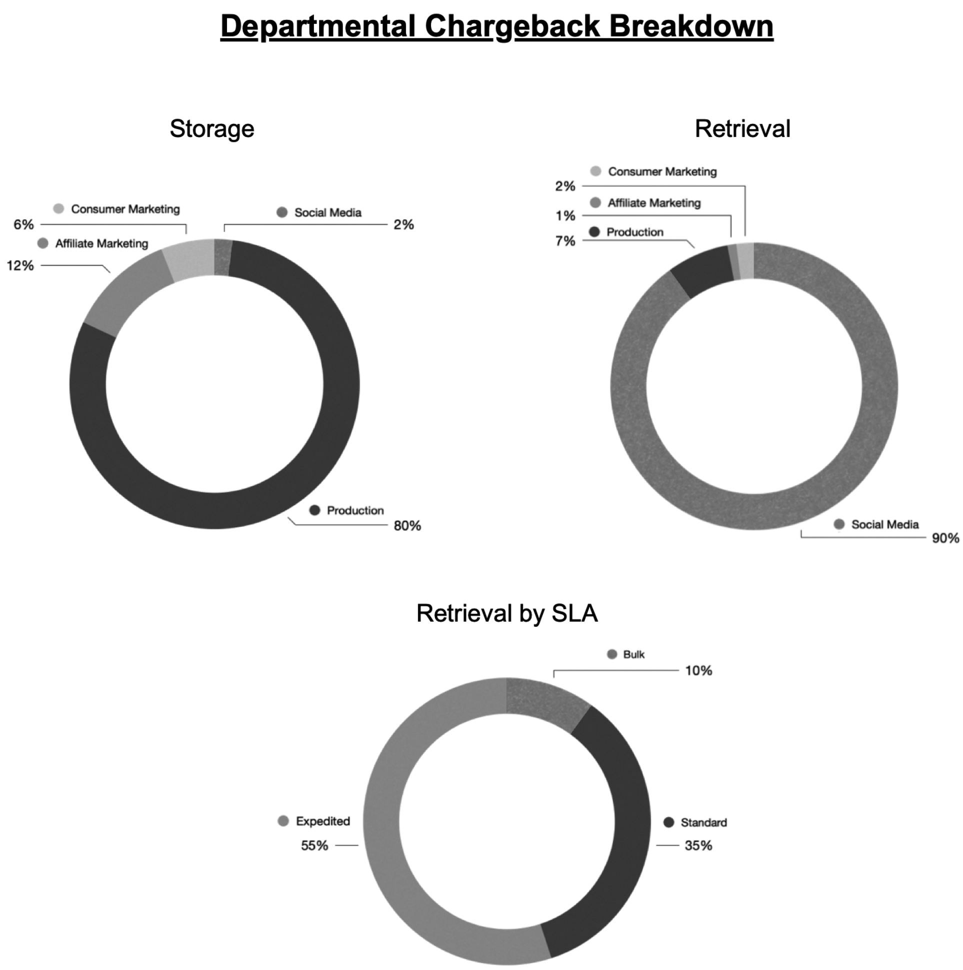

Figure 2: Departmental chargeback breakdown

storage and egress of various project assets are obscured into a nebulous bill. With correlation, the administrator can associate storage and egress costs to project and job codes. This facility has worked on 22 feature films across four studios and can now pinpoint how much each project cost, enabling more granular tracking of profitability across jobs. Figure 3 illustrates a sample breakdown.

- *Lifetime cost of assets:* Throughout the media pipeline, assets will flow from on-premise to cloud storage, and eventually to a glacier or cool/archive blob tier. With the MAM system, end users are empowered to perform their own self-service content restores as needed. The MAM system will update status flags showing what content is available and what content is still in the archive. When assets are brought back online or moved to a different class of cloud storage, there is a cost associated. Sure enough, those assets will also eventually need to move back into archive again. With correlation against the job history of the MAM system, one can ascertain how many times this asset has been moved, where it was moved, when it was moved, and who moved it, and at what service level (for retrieval). This knowledge informs an administrator about storage optimisation. If an asset is frequently accessed from the archive (enough so that it incurs the early penalty), those assets should be flagged

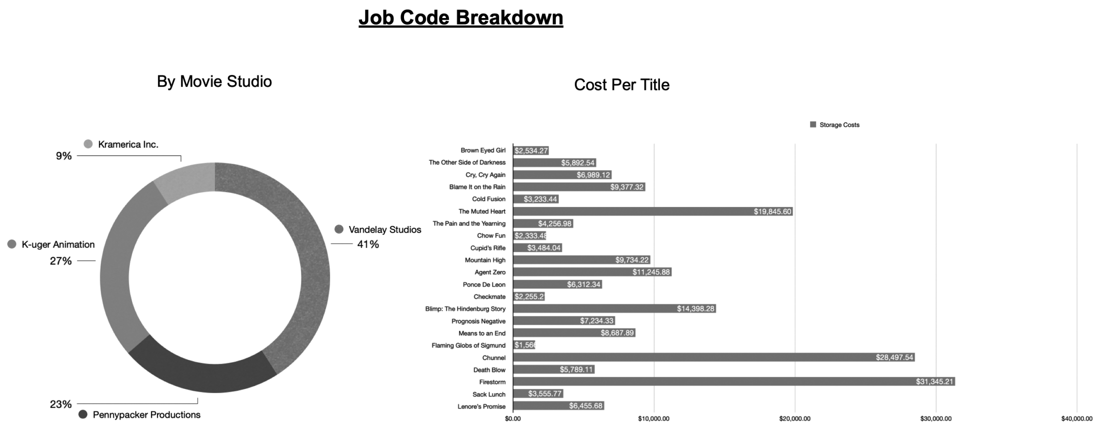

Figure 3: Job code breakdown

with a unique metadata status that keeps these commonly used evergreen assets on tier 1 storage indefinitely. Knowledge of the lifetime cost of storing assets also informs decisions around retention policy. Some companies adopt a 'keep everything all the time' policy, citing that 'storage is cheap', while other companies need to make every penny count. Understanding the cost to store assets combined with a 'chargeback' model puts the hard decision back on the content custodians to either keep paying to store assets on the current storage class, move them to the lowest/cheapest tier, or to purge them altogether. Table 3 illustrates a sample specific to individual assets, but the same reporting could be applied to entire collections of assets. Leveraging the MAM system's metadata about assets and projects, a report could be built to correlate accordingly.

THE NOT-TOO-DISTANT FUTURE

Expanding the scope of 'time and cost' beyond storage to include compute resources such as cloud transcoding and AI/ML will enable companies to make more informed decisions of where, when and how they process media. Workflows could eventually incorporate the variables necessary to understand the source asset's current storage location, intended transcoding profile and delivery target, and make contextual decisions on where and how best to process and deliver the media accordingly. These are the types of workflows that will be necessary to bridge the gap between now

Table 3: Lifetime cost sample report

Project	Asset type	File size (GB)	Date created	Dates archived	Dates retrieved	Lifetime cost on cloud (including egress)
Chunnel	DCP	154	14/1/18	14/4/18	31/10/20	US$58.18
				11/1/20		
Firestorm	AXF	172	22/7/16	22/10/16	3/8/19	US$128.22
				15/8/19	13/2/20	
				1/3/20		
Ponce De Leon	MXF	84	16/2/20	16/5/20	N/A	US$9.80
Prognosis Negative	IMF	96	30/4/17	22/10/20	N/A	US$76.26

and the future, as outlined in the MovieLabs 'Evolution of media creation' white paper.[9]

CONCLUSION

Having a well curated MAM system with relevant technical, descriptive and temporal metadata makes users happy and will drive adoption of the system. Having to worry about where assets are stored is not something users want to think about. A MAM system's proxy workflow clears this first hurdle. Providing time and cost metrics to users and administrators helps to better manage expectations around archive and restore workflows from the cloud.

Correlating MAM metadata with cloud storage metrics allows companies to understand cost implications for the long-term archival storage of assets. Total cost of ownership and chargeback mechanisms can help justify the expense of a MAM system, as well as the underlying infrastructure to support it, both on-premise and in the cloud. Cloud pricing has many variables; understanding how they can benefit the company will aid in storage cost-optimisation. Correlating cloud pricing variables to users, departments and job codes has equal benefit.

As companies look to adopt or expand an existing MAM system to incorporate cloud archive, they should consider the user experience and what intelligence can be gleaned by adding time and cost variables from the cloud service provider into workflows and downstream reports.

References

1. MovieLabs (2019) 'The evolution of media creation', available at: https://movielabs.com/production-technology/ (accessed 23rd August, 2019).
2. Kaufman, D. (2020) 'Lost Lederhosen showcases virtual production at HPA Tech Retreat supersession', available at: https://hpaonline.com/lost-lederhosen-showcases-virtual-production-at-hpa-tech-retreat-supersession/ (accessed 27th February, 2020).
3. Amazon (2021) 'Amazon S3 Pricing', available at: https://aws.amazon.com/s3/pricing/ (accessed 19th November, 2020).
4. Amazon (2020) Amazon S3 Reduced Redundancy Storage (RRS) https://aws.amazon.com/blogs/aws/new-amazon-s3-reduced-redundancy-storage-rrs/ (accessed 19th November, 2020).
5. Amazon (2020) 'Availability Zones', available at: https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html#concepts-availability-zones (accessed 19th November, 2020).
6. Amazon (2018) 'Tagging Best Practices: Implement an Effective AWS Resource Tagging Strategy', available at: https://d1.awsstatic.com/whitepapers/aws-tagging-best-practices.pdf (accessed 19th November, 2020).
7. Josyula, V., Orr, M. and Page, G. (2011) 'Service lifecycle management', in 'Cloud Computing', Cisco Press, Hoboken, NJ, pp. 256–261.
8. Cella (2019) '2020 In-House Creative Services Industry Report', available at: https://www.cellainc.com/insights/in-house-industry-report/ (accessed 19th November, 2020).
9. MovieLabs, ref. 1 above

