THE BROADCAST BRIDGE
Connecting IT to Broadcast

Productive Cloud Workflows

Essential Guide

EG

ESSENTIAL GUIDES

Introduction

By Tony Orme, Editor at The Broadcast Bridge

It's quite incredible to think that real time broadcast signal workflows are now actively encouraging software processing. It wasn't so long ago that video images had to be processed in hardware to meet the tight timing constraints that live video processing demands.

With advances in seemingly unrelated industries such as finance and medical, huge investments, far in excess of anything broadcasting could justify, have delivered datacenters that provide the infrastructure needed to build SaaS systems capable of processing and shifting real-time live video and audio. Although these new infrastructures have provided us with incredible opportunity, we must now think more in terms of IT to make the best of these new systems.

Rip-and-replace is at the core of cloud and virtualized SaaS systems. Although we are not actually removing or replacing any hardware, we are creating and deleting instances of entire servers with a mouse click, or an automated API. This reminds us that datacenters, whether on-prem or in the public cloud, are not infinite in size. It's just that the amount of resource they provide is flexible and huge, especially when we consider the public cloud offerings. Again, as an industry we are riding on the crest of the innovation from other industries, and this only plays to our advantage.

One of the key aspects of moving to virtualized and cloud processing is that we must think really hard about our workflows. Merely taking static workflows and pushing them into the cloud will not deliver the efficiencies and resilience that cloud computing offers. Therefore, we must use this migration as an opportunity to analyze our workflows to the nth degree. We will then be able to find where we can apply the dynamic principles of rip-and-replace that makes cloud computing the incredible opportunity it currently is.

Collaboration is proving its worth and the application of APIs delivers orchestration and interoperability that could only have been dreamed of a few years ago. RESTful APIs with JSON data exchange has opened new methods of working that release broadcasters from the constraints of static custom interface protocols to deliver flexibility and scalability that delivers the orchestration collaborative systems demand. This further develops scalability and provides an element of future proofing so that systems can be easily expanded later as new solutions appear.

One of the major changes that many broadcast professionals find themselves addressing is how to fathom ingress and egress overheads as costs and latencies are often cited as reasons to resist adoption. However, the solution is surprisingly straightforward, and that is to keep all the signal processing, storage, and delivery in the cloud. Once we've streamed the media into the cloud, the next place it should come down to earth is at the viewers device, whether it's a smart TV or mobile device. Advanced methods of object storage for both near and off-line facilitate all the storage and security options any broadcast facility should ever need.

Tony Orme.

Cloud and virtualized computing are providing broadcasters with the greatest opportunities seen since the 1930s when the first television pictures were transmitted. The flexibility, scalability and resilience with which public could providers operate is far in excess of the systems any broadcaster could ever provide with their own on-prem design.

Tony Orme
Editor, The Broadcast Bridge

Productive Cloud Workflows

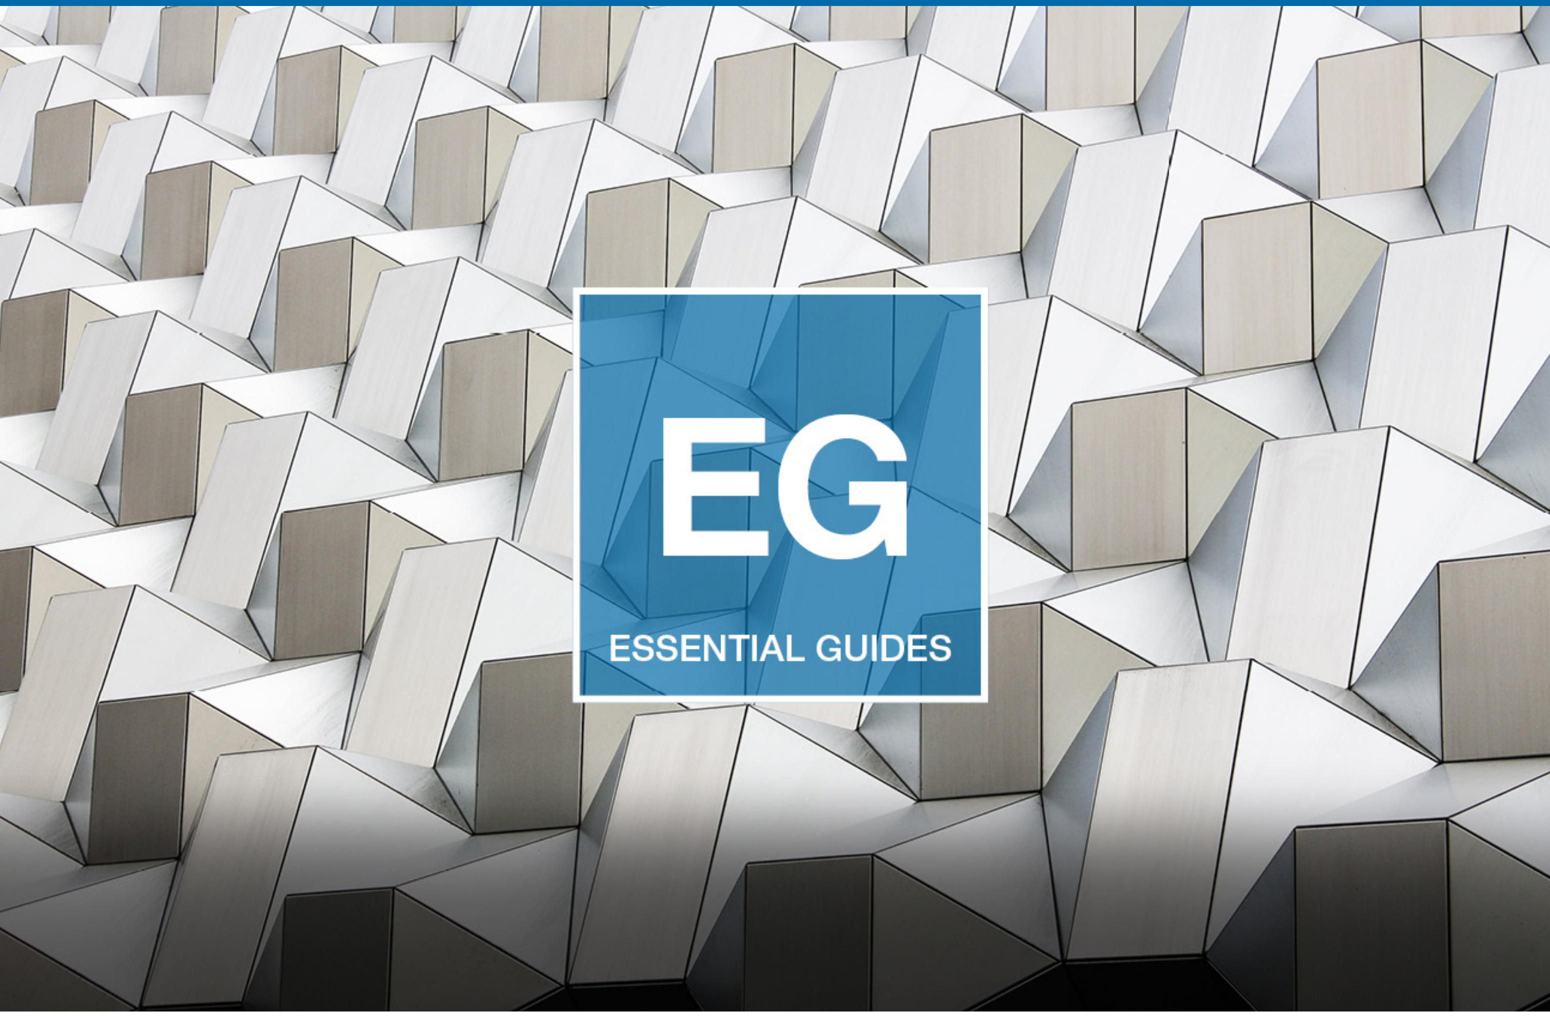

By Tony Orme, Editor at The Broadcast Bridge

IP is an enabling technology that facilitates the use of data centers and cloud technology to power media workflows. The speed with which COTS (Commercial Off The Shelf) hardware can now process data means video and audio signals can be transcoded, edited, and transferred with speeds that are fast enough for real-time live and file-based workflows.

However, there is much more to building cloud-based broadcast facilities than just running microservices and SaaS applications. Software systems must be seamlessly integrated to provide stability and dynamic response to viewer demand.

A combination of RESTful APIs and JSON data objects are needed to maintain compatibility between processes that allows them to transfer and process video and audio signals.

Trying to replicate baseband and on-prem workflows directly in the cloud is fraught with many challenges. These include overcoming latency and completely changing our working practices to agile methodologies, including rip-and-replace. The dynamic nature of cloud systems demands that workflows must be designed with effective strategies to make the most of the scalability, flexibility, and resilience that cloud promises.

Although we can take existing workflows and replicate them in the cloud, to truly benefit from the many opportunities and efficiencies that cloud can offer, we should always design systems that are cloud aware.

Scaling Sideways

As a broadcasting community, we have spent the best part of eighty years committed to hardware technology. This is mainly due to the speed with which video must be processed, especially in real time. The sampled nature of video relentlessly delivers frames of video between 17mS or 20mS apart, and as the size of the raster increases then so does the data rate. Moving from SD to HD created five times increase in data rate and moving from HD to 4K quadrupled this further. Consequently, video processing has always been at the forefront of technology, and until recently, only hardware processing has been able to deal with these colossal data rates.

Developments in seemingly unrelated industries such as finance and medical have seen massive investment in COTS style processing, leading to the development of software defined networks and data processing. Finance with micro-trading, and medical with machine learning, has led to high volume data processing with incredibly low latency, exactly the type of processing that broadcasters demand for real-time live video and audio processing.

The development of the datacenter has further led to the concept of virtualization. Modern operating systems create events for input/output (I/O) peripherals such as keyboards, mice, and ethernet interfaces. This removes the need to for higher level applications to constantly poll the I/O devices leaving much more time for the CPUs to process applications.

Virtualization takes advantage of this as the access to all the I/O devices can be divided across multiple operating systems sharing the same hardware, thus improving the efficiency with which servers operate. It is possible to run multiple versions of an application on a single server, but the concept of virtualization creates greater separation between processes to improve reliability as well as security.

Virtualization effectively abstracts the hardware away from the operating systems that it supports, and it is this abstraction that has delivered cloud computing. Another way of thinking about cloud computing is through the idea of software defined hardware as the virtualizing software is effectively a server management system. This provides a user control method that allows the creation and deletion of virtualized servers.
The ability to create multiple operating systems on one server provides the concept of cloud computing.

These ideas lead to the ability to rip-and-replace. That is, we can create and delete servers with a few software commands or clicks of a web interface. And it is this concept that is the most exciting for broadcast engineers as it is a completely different way of working from the traditional hardware workflows. This is because we now have the ability to just delete processing engines that are no longer needed thus allowing demand led resources that can be provided to meet the peaks and troughs of the business.

When spinning up servers we are creating operating systems that are hosting the required applications. But it's important to note that even the most powerful of servers has a physical limit to the number of instances that it can host. Therefore, on-prem datacenters suffer from the same limitations as traditional hardware systems in that engineers must design for the peak production demand. Although public cloud systems also suffer from the same limitations as there are only a fixed number of servers in any one datacenter, the number of servers within public clouds is orders of magnitude greater than any on-prem datacenter resource any broadcaster could ever hope to build.

With public cloud providers continuously adding more servers, storage, and routers to their datacenters, it's extremely unlikely that any broadcaster would ever exhaust the available cloud resource.

Key to making cloud infrastructures both resilient and efficient is understanding and accepting the concept of spinning up a resource and then deleting it when it's no longer required. Automation achieves this by detecting when a workflow requires extra resource and then spinning it up as required. One method often adopted is to use message queues.

Dynamic broadcast workflows are job driven. That is, each action is given a job number and every job that needs to be processed is placed into a queue, then a scheduling engine pulls the jobs from the queue to process them. At this point, the scheduling engine will analyze how many jobs are in the queue and calculate the approximate time that will be taken to process the queue with the available resource. If the estimated time is greater than a predetermined value, then the scheduling engine will spin up new servers. If the estimated time is lower than a predetermined value, then virtual servers will be deleted thus reducing the amount of available resource.

An example of this is Ad ingest for a broadcaster. Earlier on in the week the number of Ads sent by the post houses and advertising agencies will be relatively low. But as the week progresses and we approach Saturday evening, that is when peak demand arrives, the number of Ads being delivered will increase massively. A cloud processing system may only use two transcoding instances during the earlier part of the week, but as we approach Friday, the number of transcoding instances may be increased to ten by the scheduling engine. And when the Saturday cut-off date has passed, the extra eight transcoding servers will be deleted. It's important to note that the server instances will be actually deleted so as not to use any of the server resource. In pay-as-you-go cloud systems, this provides a huge saving for the broadcaster as only the servers needed to meet the demand will be actioned and hence, paid for.

As this whole process is effectively automated, the business owners are responsible for determining how many servers are made available to the scheduler. Therefore, the decision on available resource becomes a commercial undertaking and not necessarily an engineering one.

Software In The Cloud

Unix was designed around the idea of small programs performing specific tasks that could be pipelined together. This provided much greater flexibility and reliability than writing a new program for every task. But as computing developed, applications moved to monolithic designs that were huge. Consequently, they were difficult to maintain and team collaboration proved challenging. New releases were infrequent and extensive testing was difficult due to the extensive number of combinations of inputs and data stored.

To overcome these issues, modern software applications are moving back to smaller, self-contained programs called micro services. These are much easier to maintain and test as they have a reduced domain of input with a much better-defined output range. Consequently, they maintain greater reliability, flexibility, and scalability.

Examples of microservices are programs that provide specific tasks such as color correction and subsampling, transcoding, and YUV to RGB conversion. The idea is that multiple microservices are daisy chained together to provide solutions to complex workflow needs. Furthermore, they excel in cloud type environments.

To help with deployments and management of microservices, containers are often used. These are a lightweight alternative to virtualization but still provide a pay-as-you-go model through orchestration. This is a separate service that automatically spins up servers and starts and stops microservices on demand. Whereas virtualization requires a hypervisor management system that acts as an interface between the CPU and IO hardware, each container runs on the server operating system and provides a contained area of operation.

Containers derived their name from the early Linux cgroups, which later became Linux Containers (LCX) and then other solutions such as Kubernetes and Docker were derived. They allow the host servers operating system to allocate a certain amount of CPU, RAM, and operating threads to each container, thus providing a guaranteed resource allocation.

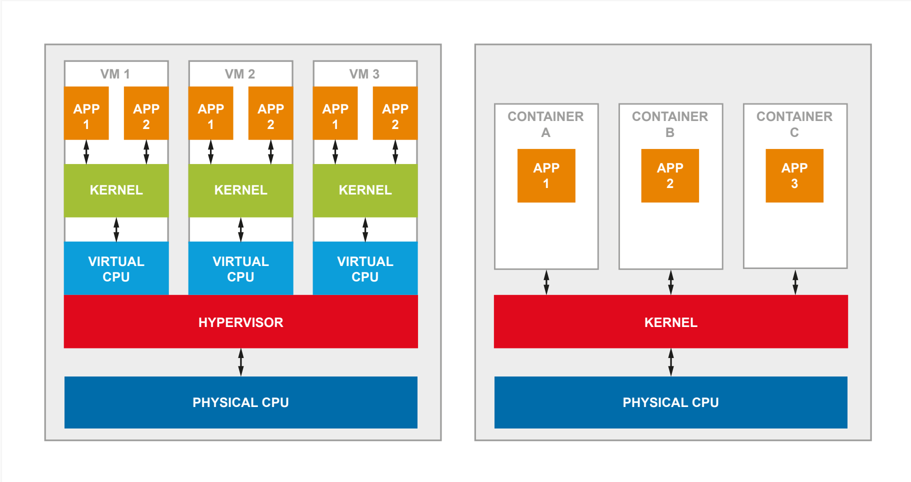

Fig 1 – the diagram on the left shows virtualized servers and the diagram on the right shows container systems. Virtualization is more flexible, but containers are much more lightweight.

Furthermore, whereas the virtualized hypervisor solution can provide multiple and different types of operating system, containers are restricted to the host server kernel operating system. However, only the operating system user mode modules are provided in the container, so they stay lightweight and provide a guaranteed operating environment with which to work from.

Communication and control for microservices or larger apps running in either virtualization or containers relies on RESTful (Representational State Transfer) API interfacing. The RESTful API is particularly useful for cloud applications as the protocol is based on HTTP (Hypertext Transfer Protocol), that is, it uses commands such as GET and POST found in web server applications. This makes integration into cloud systems much easier as the cloud service providers use the RESTful APIs as the basis for public communication with their servers.

Cloud computing relies on the API being well defined, and in the modern agile computing paradigm, this is well understood and supported. The APIs are scalable and easily maintain backwards compatibility so that upgrades are quickly facilitated. The stateless operation of REST requires the client to server request to provide all the information needed for the server to process the data entirely. The client cannot hold any state information about the task being processed.

For example, if a client is requesting a transcode operation from a microservice it must provide the source file location, its destination, and all the parameters needed for the transcode process.

Data is exchanged between RESTful services using JSON (JavaScript Object Notification) type files. This is a lightweight and human readable text format which is programming language independent and has a familiar syntax to 'C' type languages. By human readable, we mean that JSON provides self-describing information within the file using parameter-value pairs. It's similar to XML except that it is shorter and much easier to parse.

It's important to note that REST and JSON are not intended to transfer large media files, instead, they provide file locations for the source and destination files as well as the data values needed for the operation. For example, a JSON file may contain the raster size and frame rate of the source file along with the raster size and frame rate of the destination file for an upscaling process. The file locations will be network mapped drives or storage objects similar to those used in AWS or Azure. The actual media files are accessed directly from the object storage by the application, and it relies on the cloud service provider delivering adequate network capacity and bandwidth to transfer the data for processing.

Keeping Media In The Cloud

One of the challenges broadcasters often consider when moving their workflows to cloud infrastructures is that of ingress and egress costs. Although these can be significant, it's worth remembering that keeping media files in the cloud is a much more efficient method of working than continuously moving files between the cloud and on-prem. This requires a significant change of mindset but does pay dividends for costs and security, as well as flexibility and reliability.

Cloud object storage is not only an efficient and flexible method of storage but is potentially much more secure than on-prem storage. All the major cloud service providers use a token to describe the storage object of a particular media file. In appearance, this is very different to the hierarchical file systems often found with Windows and Linux, instead, it's a unique long sequence of characters. When the media owner generates the token, they also include the type of access rights and available time. Not only is this unique, but if the media owner suspects a security breach, then they can remove the token from the media object, thus stopping access.

When a client sends a REST API command to a server service, such as a transcoder, included in the JSON file will be the token for the media object that is being accessed. Furthermore, the media owner can see who has used the token, where and when, thus providing a forensic audit trail for the high value media.

Cloud service providers also provide backup storage such as Glacier in AWS, or Archive storage in Azure. These are long term storage systems that are much cheaper than the instant access storage as they often take several days to retrieve the media asset. This is similar to LTO used in on-prem datacenters. But if all the costs are considered, including physically hosting the LTO machine, supporting and maintaining it, and finding somewhere to store all the tapes, then this type of cloud storage can often be a better alternative. Again, the media files are kept in the cloud so security is maintained and network bandwidth is supported by the host cloud service provider.

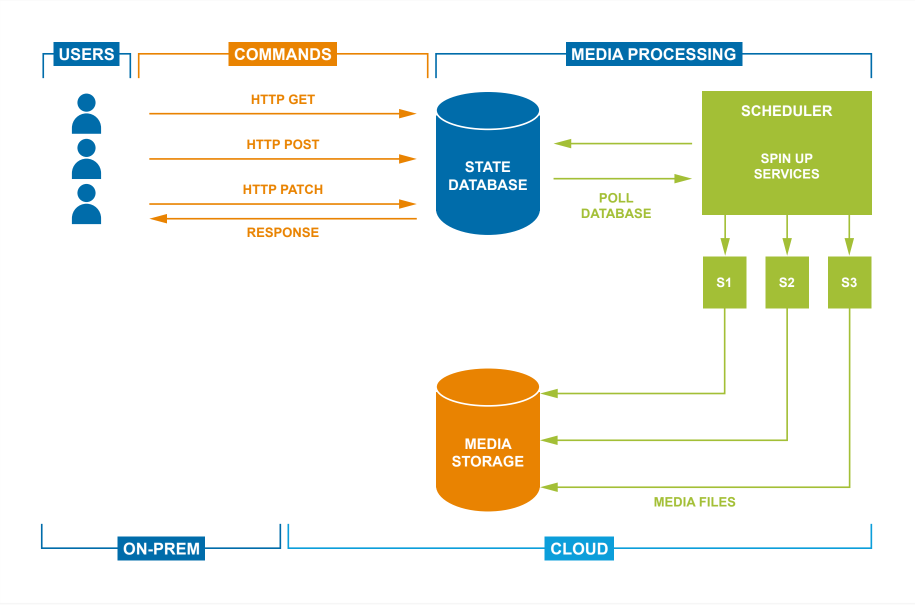

Fig 2 – the media processing is accessed through the RESTful API using HTTP. When a user has requested a service, such as transcoding, it will store the media storage information as well as the transcoding parameters in the database. The scheduler polls the database for jobs and when it receives them it will spin up new resource as needed. The processing servers will access the cloud media storage directly for both the source and destination files.

For example, a ninety-minute HD baseband media file will use approximately 537GB of storage. AWS Glacier costs approximately 0.36 cents per GB of storage per month, resulting in a cost of $1.93 per month for a ninety-minute asset, or $23 per year.

Using this example has a major advantage as it allows the business owner to make an informed decision on whether the asset is worth storing or not. This is clearly a tough call, because none of us know what the future has in store, but these decisions are put firmly back into the hands of the business owner and not the engineer. One of the challenges we have is knowing what to keep and what to delete. This provides a straight-forward metric to help establish this.

Making Cloud Work

Cloud workflows require a different type of thinking, the dynamic nature with which cloud operates demands a transient approach to workflow design. To make cloud efficient, processes that are not being used must be deleted, and this in itself requires systems to constantly analyze how the overall cloud system is responding to the operational demands.

Metadata availability is in abundance within the cloud and broadcasters can take advantage of this. The enormous amount of monitoring data available can be analyzed to understand where components within the cloud infrastructure can be closed or deleted so they are not unnecessarily using expensive resource. Furthermore, they simplify systems making them more efficient and reliable.

The Sponsors Perspective

Working In The Cloud – Productivity Is Always The Key

By Reuben Cohn, Cloud Transcoding Product Manager at Telestream.

We've encountered media companies along all aspects of migrating their workflows to the cloud. Some with large on-premises media processing capabilities are just beginning to design their path, while others have transformed some of their workflows to be cloud-native, and still, others are all in – meaning all their media files reside in the cloud and are processed there.

We also see virtual deployments that run in the cloud, but they are not cloud-native SaaS. Our customers need to be met where they are on their journey to the cloud, and we have the breadth of products and services to be able to do that.

Because of that, Telestream is the defacto choice regardless of infrastructure.

Centralized Workflows And Elastic Scale

A major advantage for moving media processing workflows to the cloud is the cost savings that results from being able to centralize them and scale horizontally and vertically as needed. The need to design media workflow architectures for peak demand is eliminated.

Shared storage is typically the first step in building out a communal media production workflow. Increased media access and collaborative possibilities positively impact creativity but also present new non-creative burdens as ingest and delivery processes become more critical. Ensuring that best practices are repeatedly followed is difficult as these non-creative processes are distributed across multiple systems and done by a variety of people. When creative people ingest and package/ deliver media or are impacted by ingest processes not properly followed, creativity and workflow efficiency is negatively impacted.

Many customers take advantage of shared storage in the cloud, but they haven't been able to maximize the benefits because they haven't realized their processes and workflows are not suited to a centralized model. Hence, we recommend to not only deploy shared storage but to also centralize workflows for cloud-based media. With remote workers and workgroups all using the exact same workflows, ingest and delivery, file naming, folder structure, workgroup rights and privileges, and other processes are normalized across the organization, and they can be scaled up and down as media processing demands fluctuate. When these cloud-native media workflows are normalized and automated, they become productive and efficient, and creative staff is freed up from mundane tasks and can create better stories.

In one example, a Telestream broadcast customer, has done exactly that. By centralizing the reception and processing of media across all their properties in the cloud, they have been able to eliminate duplication of not only workflows and media, but also much of the equipment that was needed to process this content at multiple locations across the country. As a result, they have achieved significant cost reductions, and they are able to share content within the organization more easily because it is all being processed in a central "location" in the cloud. From there, various workgroups access the content from centralized, cloud-based storage. By doing so, they have turned what used to be a laborious process of transferring files around into productive, efficient workflows. Plus, they're able to leverage dynamic scaling afforded by the cloud, and they do not need to build out for peak demand.

Updated Mindset And Process Enables Cloud-based Media Workflows

Several factors are fueling the transition to cloud-based broadcast operations and make it highly relevant to today's business landscape.

Cloud Infrastructure Adoption – Across the Media & Entertainment industry, there is a mindset change underway from owning and managing a private datacenter to adoption of cloud infrastructure for storage and management of media content.

SaaS Availability - As cloud adoption becomes more ubiquitous there is a need to have cloud native software services available to execute media workflows. Bringing the Telestream Media Framework to the cloud allows for continuity for our customers' workflows by using an underlying technology they can trust/

Elastic and Scalable – Encoding.com powered by the Telestream Media Framework can meet the expectations of rapidly changing business requests by expanding and contracting based on need at the time of execution. There is no need to overbuild or put off workflows to a later date. We can provide the service now and scale it up or down on demand.

Remote Availability – As the decentralization of workers accelerates, cloud SaaS products allow remote workforces to more easily collaborate and not rely on hardware and locally installed software.

Media Company Personnel Skills – There has been a change in who manages the operations of the media companies. More and more, there is a core competency in software development, and the Encoding.com API, for example, is the industry's most mature, well-documented, and feature-rich API for cloud encoding. Whether you're a developer or a video engineer, our API makes it easy to migrate your VOD processing to the Cloud.

Data Gravity

There is a term we use at Telestream, "data gravity," that describes how to consider where media is stored and how to work with it. Data gravity means the larger the media files the more they will attract the necessary processing to them – they will not be moved. Moving large media files around can be expensive and time-consuming. Hence, Encoding.com brings media processing to the media across almost any geography.

Hybrid Workflows

There is still a need for hybrid workflows as the transition to the cloud is realized. Even as distribution workflows are moving more and more to the cloud and the final mile of supply chains can take place entirely remotely and in an automated fashion, the origination of assets and the initial editing of raw content often takes place on the ground with some human interaction. Our software, such as the Vantage Media Processing Platform used for workflow automation and transcoding, is ideally suited to operate in a hybrid fashion where one part of the workflow or supply chain must act on content on-prem while the next steps occur fully in the cloud. The flexibility of Vantage combined with cloud allows data gravity to dictate where and how the content is processed, automating any non-creative processes both at the point of origination as well as every step along the way as the content moves to the cloud and on to its final destination.

API-First Approach

An API-First approach means that every feature is accessible through the REST API. All the functionality of Encoding.com can be programmatically controlled through a rich API from remote media workflows. Having such a complete REST API empowers broadcasters to automate, improve, and expand existing media processing workflows or test and implement new ones quickly and efficiently.

In Conclusion

A changing landscape has opened the door for many media companies to begin their journey to the cloud. While they may be at differing places on that journey, we offer a variety of deployment models that ease the transition. Whether on-premises, in virtual environments, as cloud-native services with Encoding.com, or in a hybrid combination, Telestream can meet customers where they are.

THE BROADCAST BRIDGE
Connecting IT to Broadcast

CI
CORE INSIGHTS

EG
ESSENTIAL GUIDES

MEDIA

ARTICLES

For hundreds more high quality original articles and Essential Guides like this please visit:

thebroadcastbridge.com

6/2022

Supported by
telestream