sdvi Article



Cloud Migration in Five Steps

Moving a media workflow into the cloud is no longer uncharted territory thanks to early adopters who have already demonstrated the benefits. Several of the largest media operations in the world have moved their content supply chains into the cloud, where there are no capacity barriers, bottlenecks or idle equipment. We know from their experience the model works—it saves money and time, and it's inherently adaptive.

The prevailing question is now, *how do you get there in the least disruptive way?* This guide provides a template for making that assessment. No singular route applies equally to all operations, but there are general steps that can be applied universally to create the optimal migration path for each individual operation.

We propose a five-step strategy that can be applied to all or part of the media supply chain, which we've broken into common segments or workflows for individual evaluation, to determine an optimal entry point. The common workflows are enumerated in Step 1 and elaborated on further in separate section. The five steps are as follows:



Let's start at the beginning...

1 Map the Media Supply Chain

This step involves:

- Collecting information about the supply chain
- Identifying pain points where change is most urgent

Here, we use the term "media supply chain" to refer to the full content lifecycle from receiving to distributing and archiving and all the processes in between. Existing on-premise media supply chains consist of multiple different systems, from multiple different vendors, often delicately tied to-gether, sometimes with custom software. They are inherently brittle, and frankly, once you get them humming along, you don't want to touch them. The problem is, change is imminent in the digital realm, and this on-premise conglomerative system is limited. Viewing it now as an evolutionary step toward full digital immersion into the cloud will help us define that migration path.

Start where you are. Determine what the content lifecycle looks like today. What is the track record of the operation—being asked to do more with less? Left-field requests from the non-engineering end of things? Workload feast or famine? A little fight-flight-freeze in every execution? Let's bear these conditions in mind as we envision a highlevel view of this lifecycle to map out the media supply chain.

First off, contemplate the breadth of the media supply chain you are trying to manage. Where does it start and end, and what are the individual processes being managed? (We'll consider this in a linear fashion, bearing in mind the cloud can be linear in any direction. This linear model does not lock constituents into the analog-like linear model as earlier digital work-flows configurations did.)

If you look at your supply chain as a broadly leftto-right process, what are the big steps you go through? Divide it accordingly where you find manageable segments or workflows. Once these workflows are defined, you can begin to home in on that optimal entry point for migrating the media supply chain into the cloud.

Let's use a typical media content operation to determine where the supply chain starts and ends, and what workflows exist. This is a fairly high-level view of a generic media supply chain divided in content workflows:

- Receipt
- Normalization
- Logging
- Localization
- Modification
- Distribution
- Archiving

(What each of these workflows comprise is further hashed out in Appendix A.)

Consider the distinct conditions within these workflows in terms of aging out, reaching capacity and other replacement-cycle criteria. Contrast these conditions (x) with the functions within the sub-workflow where demand is escalating most rapidly (y). The resulting intersection is likely an operational terminus or at the very least, a bottleneck. These are the pain points to factor into laying out a migration path.



2 Determine the Entry Point

This step involves:

- Envisioning an ideal media supply chain
- Drilling down into the pain points
- Identifying the primary goal

Once the workflow pain points (or the likelihood of them in the near future) are defined, focus on the circumstances around them. Who does what, in what order? How much effort is duplicated across teams or work units? Which steps are automated and which require people? How long does each step take? How many emails, spreadsheets or manual steps are involved? Identify the root cause of the pain point.

The optimal entry point will vary according to the scope, size and reach of the individual media operation. Does it accommodate 12 languages or two? Is it local, regional or global? Is it a single entry point or on every platform under the sun? Other factors might include simplicity of migration, ROI measurability, urgency of demand and others.

Assessing Optimal Entry Points

If content receipt is the primary bottleneck, it may be the optimal entry point. Perhaps the next logical step would be to migrate archiving into the cloud, which is rapidly becoming less expensive, reliably secure and more easily managed than on-premise storage.

Content distribution is another possibility, particularly if content is delivered to multiple places, with multiple disparate systems dealing with each platform. Managing that in the cloud eliminates idle capacity and provides cost-calculation accuracy. Don't leave any stone unturned in this process. Talk to everyone involved—from the operators performing the tasks to the people supporting them, including those delivering the content and those waiting for it. Once you have a comprehensive idea of how things work today, consider how you want it to work in an ideal world. What does your ideal supply chain look like in its future state?

Imagine the process in a system with no limitations and draw it out. Don't think within the constraints of your current system, but according to the needs and demands of the operation now, and in years to come. Sketch out how you want content to flow, what needs to happen at each step, and what you want the outcome to be. Don't go to deep; just form an overall idea that can be simply articulated, e.g.:

"This system needs to reduce manual intervention by x, increase capacity by y and decrease costs by z. It needs to accommodate (local, regional and/or global) content receipt parameters, automated metadata tracking, format detection and transcoding for distribution in x number of formats for x platforms, in real-time and on-demand, plus master-format archiving with instantly retrievable files through any of the number of transcoding engines."

Let this ideal be the overarching objective, but not a concrete target. Just getting started is the point, and nailing down a quick success will generate momentum. The "ideal" will likely be refined as the process unfolds. Flexible thinking is important. Set an objective that is well-reasoned but not etched in stone.



Next, define individual goals that move the operation toward the objective. Start with the operational pain points. Where are the bottlenecks? What processes are most labor-intensive? Maybe you're dealing with several far-flung content receipt sites taking in multiple tape formats and file types that must undergo some form of manual validation. Do a process evaluation to see where automation would be most beneficial. (See, "Identifying the Processes for Automation," Appendix B.)

Determining the optimal entry point may be a simple matter of what will have the biggest benefit or be the easiest to reach. Once this entry point is determined, the next step is to develop and convey the strategy for achiev-ing this primary goal.

3 Develop the Strategy

This step involves:

- Sharing the primary goal
- Forming a team to support it
- Identifying metrics around targeted results

Once you have a solid idea of your primary goal, communicate it, starting at the top. If you're this far along the path, you know by now the business case for cloud migration makes itself. If you're able to obtain an executive mandate—great. If not, you can still make the move a top priority by taking a leadership role. Dynamic leadership is consistently linked to successful cloud migrations.

With that in mind, share the goal with key stakeholders in the media supply chain. Emphasize how migrating to a cloud-native infrastructure will eliminate manual tasks and free up resources for more critical objectives. Identify the area of impact on the workload and on the operation at large. Prepare the affected community. If your migration entry point is No. 1 content receipt, "Technology is nothing. What's important is that you have a faith in people, that they're basically good and smart, and if you give them tools, they'll do wonderful things with them."

– Steve Jobs

for example, give suppliers plenty of time to accommodate a new platform in case their own business decisions are affected. Discovery did this with a roadshow for all their top content suppliers to let them know things were going to change and how to prepare.

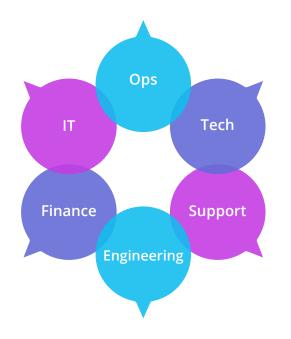
On your own turf, seek the input of those directly affected. Establish that the goal is non-negotiable, but the migration strategy is adaptive. Talk to producers, QC techs and support operators. Get their feedback. Prepare for resistance. Resistance to organizational change is a field of study all its own. Reasons range from the unknown impact on job security to a loss of exper-tise to just finding the time to learn new skills while doing the job at hand. One factor consistently mentioned in the field of change resistance is the one most easily addressed—communication.



"When you decide to make widespread changes, proper communication about why you're making the changes and how you plan to implement them is essential," writes Sampson Quain in the Houston Chronicle. "If your employees have no idea why you're asking them to change protocols that they're familiar and comfortable with, they tend to resist those changes."

Encourage people to share doubts and ask questions, so misconceptions can be curtailed early on. Have them look for the gaps in your plan and help refine it. Let them know how the move will affect the workload distribution and skill-set demands.

Begin team-building. Identify those within the organization who are on board, and enlist their support. Empower them to address misconceptions. Consider forming a steering committee. Designate a representative from each functional area, including operations, technology, support, engineering, finance and IT. Even those not directly involved in the project may contribute something unexpected, and keeping everyone informed will help avert interdepartmental snafus. You don't want a cloud migration coinciding with IT moving finance to new servers.



Start developing the project-success metrics by identifying measurable impacts, e.g., reducing manual intervention (hands-on time) by x, increasing operational capacity by y. Maybe you want to contain the cost of maintaining a physical infrastructure, or eliminate poor equipment utilization rates and optimize capital expenditure-to-usage ratios. Just consider what can be measured going into the project so you have a reasonable notion of your baselines.

Talk to the cloud-service providers who can align their services with your goal. Remember that the processes throughout each sub-workflow are offered as pay-as-you-go microservices within the cloud from a variety of vendors. This minimizes the time spend on vendor selection, takes a huge capitol risk factor off the table and minimizes vendor lock-in if the selected technology doesn't perform as expected.

Familiarize yourself with this new pay-as-yougo provisioning and billing model. Cloud-native automation enables rapidly provisioning services as they are needed. The cloud cost model is a radical departure from the facilities cost model, so it will be worth the time to get acquainted with it.

4 Implementation

This step involves:

- Creating the implementation team
- Setting the timeframe
- Testing, tweaking and launching

Once the migration strategy is nailed down, the focus can shift to implementation. It should be fairly straightforward to create an implementation team since by now you know who is most engaged. Identify—either train or hire—an in-house cloud administrator. This could be someone already in IT. Cloud admin may not be a full-time job, but it helps to have someone on staff who understands how the cloud operates.

Establish a Governance Process

- Develop system standards
- Create a reporting and evaluation process
- Establish a maintenance protocol
- Apply across all cloud implementations

The team can help catalog a few things—e.g., network bandwidth capacity, where to apply encryption and what security requirements are indicated. Along the way, have them establish an ongoing governance process to keep the strategy on track, and to continue communicating with stakeholders so problem areas can be identified and mitigated quickly. The process will involve developing technical, connectivity, resiliency and failover

standards, plus the practices for maintaining those standards. Investing the time to es-tablish a governance process early on will pay off as migration continues because it can be adapted across all phases of a migration and beyond.

At this point, you may want to sketch out datasourcing protocols for capturing metrics. Data has its own data in the cloud, so it's helpful to think about just how granular you want to go. Having access to the cost of every individual process on each piece of content doesn't mean this data has to be continually compiled. Don't base these protocols on the available data, but rather the data necessary for establishing metrics and achieving the desired outcome.

Establish a realistic timeframe with regular progress reports due from the implementation team. When it comes to setting a date to take the new system live, shoot for the least busy time of year—if you have one. Account for holidays and other possible disruptions, including those surprise mandatory workplace training sessions triggered by some new regulation.

This is where having thoroughly communicated the migration plan across all departments including human resources—comes in handy. The move may not have direct impact on every department, those departments may have a direct impact on the move.

Back up existing servers and data if that process isn't already taking place automatically. Test connections, individual components and the system in total to confirm reliability. Make sure content is presenting correctly and the environment is secure. Check for any issues with end users so the appropriate tweaks can be made.



When you're ready to go live, route content through existing and new processes and observe the difference. Note where automation eliminates the need for manual intervention. This dual-workflow phase allows for fine-tuning while de-risking the initial migration, plus it demonstrates proof-ofconcept internally and creates a reference system for the continued roll-out.

Keep communication lines open. Expect to talk a lot, particularly in organizations where departments that haven't worked together before. Then commence the migration.



This step involves:

- Applying metrics
- Assessing the outcome
- Adapting the process for the next round

Once the migration is complete, revisit to the primary goal set in Step 2. Apply the key metrics developed in Step 3 with the data identified in Step 4.

Get post-migration feedback. E.g., how long does it now take to process one hour of content from beginning to end? How many manual interventions are now required? What is the overall cost of my supply chain versus the pre-migration model? How much more content can the team handle now compared to before? Is their time now focused on more critical or creative endeavors? What is the cost of change?

The cloud-native ecosystem allows you to constantly evolve and adapt as the business requirements change. This unprecedented visibility into near real-time processing costs will enable data-informed business decisions.

The operation can more opportunistically take additional volume deals where they make sense rather than whether or not you have the capacity to deliver them. This ability to get to revenue faster, with more overall capacity and no down time fast-tracks your ROI.

All of this information, along with the experience of navigating a full migration cycle, will form the groundwork for the next workflow migration—or full immersion where it makes sense. At this point, you can identify the methods and practices that worked well and determine how to apply them to other areas of the media supply chain. In rolling out cloud-native content receipt, for example, you may discover a technique that would have a transformational impact on content delivery.

That's part of the beauty of this migration model—it's iterative. Once you have migrated your initial workflow, you can modify and adapt the methods you used to make the next one even more efficient and streamlined. This is where the governance process will come into play and also where it will continue to be refined.

So now return to Step 1, identify the next area of the workflow for migration, and go from there.



Appendix A.

The Common Workflows

Determining which supply chain workflow is most suited to cloud migration:

- 1. Receipt
- 2. Normalization
- 3. Logging
- 4. Localization
- 5. Modification
- 6. Distribution
- 7. Archiving



For our purposes, the media supply chain starts with receiving a piece of content from a producer. That piece of content may have existed before this point—as a record, some metadata or rights—but the media supply chain begins where you take it on board, and just accepting content from producers can be a painfully manual task in a traditional media supply chain.

Once content is received it needs to be checked for validity. Does it meet expectations and technical specifications? Is the metadata complete? Does it include all the component parts needed such as languages, subtitles and so forth? Does it pass technical quality tests? Assuming it conforms to all these required conditions, it may still have to be transcoded to a house format, depending on the incoming formats you're willing to accept.

There are a lot of steps involved in just receiving, acknowledging and accepting content into a facility. Depending on the scale of the operation and how many independent producers contribute content, this could be a huge, labor intensive exercise. How efficient is it today? How many people are involved? What are the failure rates of different providers? Do you even know? How long does it take to accept each item? For the vast majority of people, these are often unanswerable questions because there's no real way to measure or get analytic insight in to the processes.

If we go back to the left-to-right model described in Step 1. of "Cloud Migration in Five Steps," content receipt is a logical place to start. If your current mode of operation includes manual QC and having people watch content in real time, for example, starting the migration with content receipt can significantly increase opportunities to automate as well as eliminating many of the the complexities around it—including format, quality and metadata validation. Remember the evaluation guide for "Identifying the Processes for Automation" in Appendix B.

It's also a logical starting point if the objective from the get-go is a completely cloud-native media supply chain. A migration that begins with content receipt can be achieved using an

online upload portal to the cloud of choice, where files can be inspected, corrected, accepted or rejected automatically according to specification. Content suppliers can be onboarded over time while the next logical step is being executed.

All of the typical complications and obstacles are eliminated in a cloud-native content-receipt environment. Not only are manual processes automated, but automated in a way in which they only incur cost when active. Plus there is full access to analytical data, any way you want to slice it.

1ST PLACE:

WITH MOST MEDIA SUPPLY CHAINS, CONTENT RECEIPT IS THE MOST LOGICAL PLACE TO START MIGRATING TO A CLOUD-NATIVE ENVIRONMENT.

{ကြွဲ^{ခြာ} Normalization

Many media organizations have a preferred house format for their master media—something that's undergone serious deliberation around quality, efficiency and interoperability with other systems. It's unlikely, however, that all content arrives in the preferred house format. Most media operations receive content in several formats from many different places. Even with strictly enforced content delivery specifications in place, non-compliant content will still be submitted, and when push comes to shove, it will be accepted if deadlines are tight and there is no other option.

All this leads to the need to "normalize" content to the preferred house format—to make everything the same and restore order. In a traditional media supply chain, this means selecting a transcode vendor and making sure their technology can deal with all the incoming and outgoing formats—those you deal with now and possibly in the near future. Yet another new producer is right around the corner. If the existing transcoding technology can't handle it, another is added. And another. Before you know it, you have one of each, all underutilized.

Within a cloud-native environment, transcoding and other tools can be selected from any of several vendors on a per-usage basis. Welcome to the world of utility transcoding.

If you already receive your content in the cloud, this is a great place to start your media supply chain migration. No more picking transcode vendors one at a time and being limited by scale. Just pick the one you want to use for each content source and automatically route content through it, making everything the same. Normalizing in the cloud also creates the beginning of a cloud-based archive.

3RD PLACE:

WE CONSIDER NORMALIZATION THE THIRD MOST PRACTICAL PLACE TO BEGIN A CLOUD MIGRATION.

Logging

The primary goal of content logging is to enrich the content metadata to make it more searchable, more organized, and easier to monetize at any point in the future. This could include production metadata; process metadata delivered by the tools that have operated on the content; descrip-tive metadata added by operators; or AI data, including things like object tags, face detection, transcription etc.

While content logging is inherently easier to perform in a cloud environment because of the broad availability of different tools, it's not often one would start a media supply chain migration here unless the goal was to create a rich cloud-based archive.

TIE FOR 5TH PLACE:

NOT REALLY A LOGICAL STARTING PLACE IN THAT YOU'LL BE MOVING CONTENT INTO AND OUT OF THE CLOUD, THOUGH IT MAY BE USEFUL WITH WIDELY AND DIVERSELY DISTRIBUTED CONTENT. HOWEVER, ONCE YOUR CONTENT IS CLOUD RESIDENT, THIS IS AN EASY ADD-ON TO PROVIDE VALUE TO YOUR CONTENT AND PROCESSES.

$\mathbf{X} \leftarrow \mathbf{Localization}$

Anyone distributing content globally is familiar with the localization problem—sometimes referred to as "compliance." Every region that content goes to likely has its own regulatory restrictions, meaning it must be modified multiple times for each destination. Adding languages and captions is the small part of the process, but what about removing nudity, violence, objectionable language or other material not permitted for some locations?

Public cloud-based artificial intelligence services already offer a lot of value in this area, as do specialized AI tools developed for the media industry. Having computers detect and identify these areas of concern and then guide operators through the compliance process can save a great deal of manual processing time.

Plus, if all your content is already in the cloud, taking advantage of these services is as simple as adding an option to your media supply chain. We've even worked with customers who have specifically moved their content to the cloud to analyze it using AI services in order to increase the reach and marketability of their archive.

TIE FOR 5TH PLACE:

NOT REALLY A LOGICAL STARTING PLACE IN THAT YOU'LL BE MOVING CONTENT INTO AND OUT OF THE CLOUD, BUT MAY BE USEFUL WITH WIDELY AND DIVERSELY DISTRIBUTED CONTENT. HERE AGAIN, ONCE YOUR CONTENT IS THERE, ADDING THE CA-PABILITY IS BOTH SIMPLE AND VALUABLE.

Modification

Other content modification processes may be necessary before distribution. Whether that's removing ad breaks for VOD platforms, adding ad breaks for linear distribution, trimming durations for specific purposes or even creating promos specific to each platform—content will often need to be modified.

In a cloud-based environment, avoiding egress is avoiding additional cost. If you need to modify content that's already in the cloud, there are ways to do in without taking it out of the cloud. This may involve leveraging technology such as PC-over-IP, remote or proxy-based editing or other cloud-native tools.

It's important to make sure these steps are tracked within the automated media supply chain, particularly manual interventions or modifications. While modification is perhaps not the ideal place to start your cloud journey, it's a critical and simple addition once your media is in the cloud.

4TH PLACE:

NOT REALLY A LOGICAL STARTING PLACE IN THAT YOU'LL BE MOVING CONTENT INTO AND OUT OF THE CLOUD, BUT MAY BE USEFUL WITH WIDELY AND DIVERSELY DISTRIBUTED CONTENT.



Distribution used to mean getting program and ad content ready for a playout automation system feeding linear channels. Now, it can mean supporting the distribution of content to as many as 200 disparate platforms that may be linear, on-demand, electronic sell-through, over-the-top, mobile, international or even 4K with a combination of languages and closed captioning and audio formats. It's a lost revenue opportunity when a media distribution system can't modify content sufficiently to meet the requirements of a new platform.

Traditional on-premise workflow systems often fall short here because of their fixed nature. Accommodating new and emerging platforms can be an exercise in retrofitting the workflow again and again. Even so, these systems still can't scale quickly enough to meet the inherent fluctuations in supply and demand.

Cloud-based media supply chains can offer significant advantages for distribution—particularly multiplatform distribution—in terms of being able to scale the output up and down as you need it, the elasticity to try new configurations and the cost-predictability to do so. Distributing media from the cloud allows you to try out individual tools and services on demand, providing unprecedented flexibility to adapt as new distribution platforms emerge.



2ND PLACE:

WE RANK DISTRIBUTION AS THE SECOND MOST PRACTICAL PLACE TO START, AFTER CONTENT RECEIPT. IT HAS THE BIGGEST PAYOFF, AND YOU COULD DO EVERYTHING UP TO THIS POINT ON-PREM, MOVE MASTER CONTENT TO THE CLOUD PRIOR TO DISTRIBUTION, AND PUBLISH FROM THERE.

Archiving

The cloud is a universal access point for retrieval. Once content is in the cloud—and once it is analyzed, logged and stored—retrieving stored content for future opportunities becomes as simple as running a supply chain on it.

Once you are working with content in the cloud, creating a cloud archive is a practical and relatively simple move. With deep cloud storage starting from as little as \$1 per terabyte per month, its impossible to beat this price with an on-premise archive.

If you are undertaking any of the previous workflows in the cloud, adding an archive step is a simple and painless thing to do. And if you want to start with the archive, all you have to do is work out how to get the content from its existing location to your chosen cloud vendor, and there are multiple solutions to help with that.

HONORABLE MENTION:

STARTING A MEDIA SUPPLY CHAIN CLOUD MIGRATION WITH ARCHIVING MAY BE PREFERABLE WHEN MOST OF THE CONTENT IS LONG-TAIL AND HAS AN AFTERLIFE ON MULTIPLE PLATFORMS, OR WHERE THE EXISTING ON-PREMISE ARCHIVE SOLUTION IS IN A REPLACEMENT CYCLE. EVEN WHEN THE MIGRATION IS STARTED IN ANOTHER AREA OF THE SUPPLY CHAIN, ARCHIVING INTO THE CLOUD WILL BECOME MORE AND MORE PRACTICAL AS THE MIGRATION PROGRESSES.

Appendix B

Determining Where to Automate

The following list illustrates a few things to consider when determining what areas of the media supply chain would benefit from automation:

- Tasks requiring a significant amount of manual intervention, or hands-on time to complete, e.g., more than _____ hours a ____.Processes that are frequently activated and spun down.
- Those processes requiring _____ or more steps, or distinct actions, e.g., filling in metadata fields on multiple dropdown menus.
- Functions involving multiple file types.
- Functions drawing on data from several legacy systems.
- Areas where the occurrence of errors may negatively affect compliance.
- Processes requiring tracking information across multiple iterations and modifications.
- Tasks that primarily rely on the skill sets of one or just a few specifically trained individuals.
- Areas where documentation must always be current and transparent.
- Tasks with a high risk of human error due to complexity or the number of steps required.
- Processes executed by opaque legacy software, or Frankencode developed as a workaround.
- Workflows dealing with highly changeable file formats or file requirements.

For further information, visit www.sdvi.com, or email us at info@sdvi.com