

Cloud-Native Media Supply Chains in Plain English

Imagine if you could say, “Alexa, process this deluge of content for every display type known to modern humankind,” and Alexa says, “Sure. Would you like it interpolated for virtual reality as well?”

Imagine, too, that this Alexa is a portal that can be anywhere, at any time, and you can see all the processes taking place. What’s more, you only pay for a given process when it occurs—not the hardware or the software doing it—just the process itself, and only when it’s active. This is the model of a cloud-native media supply chain. Compared to legacy media asset management systems, it’s the difference between keeping a fueled truck in a garage to pick up your next delivery— or having Alexa make it show up at your door.

After all, media operations are supply chain operations—raw materials being processed into final products for consumption. The objective of a supply chain, and the objective in today’s competitive M&E landscape, is to align supply with demand, at all times, as efficiently as possible. Still, media supply chains present complications unique to the product: content.

Fox, for example, had to process all 218 episodes of “X-Files” for Hulu on very short notice. The demand exceeded their system capacity. It wouldn’t scale, and outsourcing was costly. So Fox went cloud-native and Hulu got 11 seasons of “X-Files.” Discovery’s supply challenge involved growing demand from disparate VOD platforms across Europe they simply could not meet with traditional means. Being cloud native meant they never have to leave money on the table.

On the supply side, the digital transition led to a proliferation of file and metadata formats that evolved out of vendor competition, in a rapidly changing marketplace, where standards chased technology. This left media engineers cobbling together walled gardens and scrambling for a glue to hold them together. These architectures have proved inadequate for a marketplace defined by perpetual impermanence. Changing distribution platforms and rapidly shifting consumer behavior mean demand is a moving target.

Fortunately, the Moore’s Law of network technology has intersected with the bandwidth-dense needs of media, so the supply chain can now be entirely virtualized and cloud-centric. Media businesses today can have a fully adaptable, immediately responsive and transparent supply chain.

Transparency here refers to full visibility into each individual step in the supply chain, in terms of the time and cost required to execute. This radically shifts how budget estimates are made—i.e., on a process-specific basis versus the infrastructure necessary to implement that process. It also provides the ability to ramp up each process as needed, so there’s never lost opportunity or idled capacity making the meter hum.

Beyond COTS

Larger media operations are adopting this cloud-native model for its efficiency and flexibility. Think of it: commoditized processing versus the cost of hardware, storage, networking, licensing,

support and upkeep of assets with a 30 percent utilization rate. With the former, revenues rise as costs diminish—at up to 25 percent in one case where dedicated edit suites were swapped for standardized workstations, and not just COTS hardware.

“The objective of a supply chain, and the objective in today’s competitive M&E landscape, is to align supply with demand, at all times, as efficiently as possible.”

This standardized workstation is more of a generic portal through which individual processes are activated and monitored, while underlying minutia are automated. There is no manual metadata entry. No one has to view an entire program in real time to identify and diagnose problems or bounce it back to a content provider for reformatting. It’s like having a double line of defense for metadata management, compliance and quality assurance. One that’s cheaper, faster and inherently responsive to fluctuating demand.

So let’s take a closer look at this concept of the “media supply chain,” which can mean different things outside of an actual media operation. Some use the phrase to refer to finished files interacting with business systems. Advertisers use it from an ad-buying perspective. Here, we use it at the content preparation and distribution level, where quality control, validation, transcoding and packaging take place — from a cloud-native perspective.

Beyond MAM

What is a cloud-native perspective? It’s a way of building a workflow from within the cloud, rather than for the cloud.

Let’s look at how a cloud-native media supply chain compares to a traditional media asset

management system, starting with the MAM. Legacy MAMs were designed before the cloud was a thing. Most large-scale enterprise MAMs were either networked silos of self-contained software applications, or the protocol-linked vendor mélange of a service-oriented architecture, running on a collection of networked infrastructure. System capability and capacity is governed by the weakest link which, in all probability, reached obsolescence long before full depreciation.

These Frankensystems — ingenious as they were at the time — were created for fixed, on-premise deployment, and necessarily comprised disparate components working in tandem. They incorporated operational redundancy where processing was separate from process because of the internal silos — like a car having four separate engines for each wheel, and two more idle ones to be sure the four-wheel drive you just bought for all the ski trips yet to come is fail-safed.

There was nothing wrong with these systems. On the contrary, they represented extraordinary resourcefulness with what was available at the time. Plus, they served a purpose: moving media operations into the early digital domain. That domain, however, was characterized by dial-up internet measured in kilobits per second. Less than 20 years later, we’re watching 4K video on Netflix, with multiple other video-enabled websites open on various devices at, on average, according to Speedtest.net, 96.25 megabits per second. Meanwhile, cellular carriers are promising 5G wireless network speeds of 20 gigabits per second, and 8K video is loose in the wild.

How does that play out? No one yet knows, but it’s clear that legacy media asset management systems are conceptually out of date. Compared to a legacy MAM, and even a cloud-enabled MAM, a cloud-native system is different because it is designed within this expanding internet environment, based on unique attributes and benefits of that environment. It is inherently part of the digital architecture.

This allows software applications to be broken down into processes, or “microservices,” that can be priced based on usage. Additionally, it enables continual updating with no user interruptions. No more bringing down the entire system to try out new software that may or may not play nicely with it.

Cloud-Native Capabilities

There’s an unprecedented level of elasticity, flexibility, predictability, utilization and intrinsic failover capability with a cloud-native system. Let’s consider each of these individually.

Elasticity: As we’ve established, our cloud-native media supply chain is built from within the cloud as a virtual iteration of a physical system in a given moment. This essentially means it can change from one moment to the next and that, unlike the physical system, our cloud-native system can shrink and grow as needed. Elasticity replaces the capacity limitations of fixed systems because the cloud-native architecture can be ramped up or scaled down according to patterns of demand.

“Media businesses today can have a fully adaptable, immediately responsive and transparent supply chain.”

Flexibility: The cloud-native environment also provides the ability to go to plan B without disruption, with none of the contractual and licensing nightmares of traditional software swaps. What’s more, because the cloud-native platform is more function-focused than industry-specific, the field of developers and vendors is vast, and use cases abound. So when an entirely new capability is called for, applicable options may already exist in other supply chain models.

Predictability: A cloud-native architecture is predictable. Function drives software development, rather than vice-versa. Software development is continual rather than batch-based. It is OS agnostic. The cloud-native environment is independent from the operating system framework and all the roadblocks and barriers that entails, plus it entirely eliminates any need for over-provisioning. There’s virtually no guessing about how much a process or service will cost. According to a 2016 Deloitte Insights report, “High-speed connectivity, low-cost sensors, and cloud computing make it possible to economically allocate, track, and charge for smaller units of products and services.”¹

Utilization: The deconstruction of applications into microservices enables 100 percent utilization at all times. No asset is left idle. Pricing is based solely on consumption of services. Provisioning takes place automatically within the cloud, which means the sky is the limit in terms of available capacity. Elasticity is inherent, rather than contingent on multiple compute and license systems built for peak operations. You have a bike courier or a bullet train at your beck and call, by the hour, at an upfront rate.

Intrinsic failover: Cloud-native architectures also have no real physical boundaries. They are, by nature, multi-AZ architectures, for “multiple (geographic) availability zones.” This introduces built-in backup redundancy that is invisible to the client or the user. This global reach combined with usage-based pricing serves to opens barriers for nontraditional content contributors and emerging distribution platforms, plus the capacity to experiment with both, at a very low cost of entry.

Certainly some media organizations have cited reasons for retaining on-premise supply chains or hybrid models, but as the industry becomes more comfortable with the reliability and security of the cloud, the case for migration starts to look irresistible.

“What is a cloud-native perspective? It’s a way of building a workflow from within the cloud, rather than for the cloud.”

Why Cloud Native for Media & Entertainment

And why should cloud-native matter for M&E in particular? For the same reason that implementing tape-based workflows with files failed to leverage the full potential of the digital transition. Users were left with fixed infrastructures entirely unprepared for the explosion of formats, versioning and platform types that we now recognize as an ongoing evolution with no fully fixed characteristics.

The lack of flexibility has left media incumbents scrambling while startups shattered old models and swept up marketshare. This led to a sort of arbitrary cultural boundary delineating “new” and “old” media, one that reduced the track record of long-established operations to an aspersion. Catching up was not so much the objective as clawing one’s way back to relevance.

This is why cloud-native systems are critically important now.

The evolution of media hasn’t plateaued. On the contrary, it continues to unfold toward virtually merging with those who consume it, and not in a “Black Mirror: Bandersnatch” type of way, but the \$1 billion esports sort of way, where the 21-to-35s live. Twitch, the live-streaming esports platform launched in 2011, has 15 million unique daily viewers who spend, on average, 95 minutes a day watching live gaming. “Monday Night Football,” by comparison, logged an average audience of 11.8 million last year.

This evolution is just beginning.

MAM-to-Cloud-Native Migration

Media technology migration has always been complicated by the always-on nature of the beast. You can’t use a loaner system for a few weeks. There has to be a pathway into the cloud-native supply chain.

Media receipt, preparation and delivery is one logical area to start, particularly with the breadth of tools already available in the cloud. Processing can be rerouted into the cloud, and back to the MAM for distribution. Validation, QC and packaging — done with preferred or existing tools, and with time-based tracking metadata — takes place in the cloud-native environment.

Some have progressed further in this migration by creating a web portal for content ingest. This inherently reduces the need for physical storage types and lays the groundwork for cloud-based archiving. Once distribution is moved to the cloud and legacy archives migrated, the MAM can be decommissioned.

Why Cloud-Native Now?

Serendipity. It’s a beautiful thing. Just as the many, ongoing industry shifts are making traditional media supply chains obsolete, the accessibility of cloud-centric tools is blossoming.

One of the most relevant tipping points came when the cost of cloud storage became cheaper than the total cost of owning a selection of on-premise storage tiers. The next thing was getting content in and out of the cloud securely and quickly. These systems are now in place and trusted by some of the largest media operations in the world. For them, making this move was imperative to meet the increasingly complicated and escalating demand for their product -- one they clearly do not expect to level off or slow down anytime soon, with good reason.

Consider the two decades since dial-up modems ruled. The media of 1998 was fundamentally then

same as the media of 1978. Movies and TV shows were delivered through broadcasting, cable, C-band and theaters.

Twenty years later, broadband delivery is supplanting those legacy distribution platforms, and right now, at the University of Toronto, researchers are applying artificial intelligence to R&D in a way that promises to shatter current limitations.² AI, plus the rise of 5G, implies a yet-to-be-imagined media landscape 10 years from now.

There is only one type of media supply chain infrastructure available today that can adapt to a perpetually unprecedented future, and that would be cloud-native.

“There’s an unprecedented level of elasticity, flexibility, predictability, utilization and intrinsic failover capability with a cloud-native system.”

Notes:

1. Hagel, John, et al. Feb 12, 2016. Align price with use: Reducing up-front barriers with usage-based pricing. Retrieved from <https://www2.deloitte.com/insights/us/en/focus/disruptive-strategy-patterns-case-studies/disruptive-strategy-usage-based-pricing.html>.
2. Rotman, David. Feb 15, 2019. AI is reinventing the way we invent. Retrieved from <https://www.technologyreview.com/s/612898/ai-is-reinventing-the-way-we-invent/>.

About SDVI

SDVI brings supply chain thinking and modern technical approaches, backed by decades of experience, to Media and Entertainment. SDVI Rally is a cloud-native media supply chain platform that deploys all the applications and infrastructure needed for dynamic, responsive media operations. With cloud, hybrid, and on-premise options, Rally supports business agility, efficiency and growth with easy workflow design, elastic scaling, clear data and cost analysis, smart automation, consumption pricing, and more. SDVI's broad Application Services Market gives organizations the freedom to choose the best tools for each job.

SDVI is a privately-held company, based in Silicon Valley, California, with offices in the US and Europe.

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