Agile Media Blueprint

Creating and monetizing content using the Internet technology platform

Richard Cartwright - CTO - Streampunk Media Ltd.
Brad Gilmer - President - Gilmer & Associates, Inc.

March 19, 2018

A discussion paper from the
Advanced Media Workflow Association

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

Introduction

This discussion paper introduces concepts along with an architectural blueprint that may be used as the basis for “cloud-fit” professional media facilities of the future. The authors hope that by putting forward the ideas contained in this paper, the AMWA can begin to take a more in-depth look at the application of cloud-fit technologies in our industry. We also hope that this paper serves as a “call to arms” for those who are like-minded, and who are in a position to collaborate together to further develop the concepts introduced here.

We anticipate that this paper will be read by a broad range of people, some of whom are quite familiar with Internet technologies, and other readers who may be professional media subject-matter experts, but who have had little exposure to computer science and the concepts behind the cloud. For that reason we suggest that readers make use of the Table of Contents, taking in the information in this paper that is most relevant to them.

The Agile Media Blueprint

The Agile Media Blueprint (AMB) is a plan that enables media companies and suppliers to build highly flexible and scalable systems that run on the same platform as the Internet. This allows media companies to create and monetize content more effectively than if they were using conventional facilities. The AMB makes use of all of the hardware, software, networking and associated components used to run world-wide huge-scale systems such as popular social-media applications. We refer to all of this equipment and software as “the platform”, but please remember that every time we use that term, we are simply referring to the technology that runs the Internet.

What problem does the Agile Media Blueprint address?

It would be nice if, when a new business opportunity presented itself, a media company could quickly spin up the appropriate infrastructure and resources to take advantage of that opportunity. But with vertically integrated facilities and monolithic software applications, it becomes difficult, perhaps impossible, to take advantage of a new opportunity, particularly if the new opportunity involves even a modest level of personalization. This can be especially true if the opportunity is not initially expected to generate revenue, and if at the same time, the opportunity requires development of new infrastructure or workflows.

Today, it is not uncommon to find at least two major technical operations in a media company; one operation dedicated to traditional linear feeds, the other dedicated to serving web content to mobile consumers. This did not pose any particular problem when the web operation was small. However, as web operations have grown, visibility has increased. An Internet outage, whether for a media company or for a major social-media application, gets noticed.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

This leaves the media company in the unenviable position of running two major operations that have similar business requirements, but that are far from identical with regard to technology and infrastructure. Furthermore, the skill sets necessary to run and support the two operations are different. To make matters worse, the increase in operating cost related to running two technical operations comes at a time when the industry is experiencing competition from companies whose stock-in-trade has been developing and deploying Internet Technology. Those technology companies are bringing the expertise they developed in online retailing, banking, and stock trading to bear on the media industry, creating new products that cannot be matched by broadcasters using traditional facilities.

Media companies need a way to streamline operations, reduce the friction involved in monetizing content, and increase their flexibility in delivering content to an ever-widening number of consumer platforms, while at the same time holding the line on costs. They also need a way to deliver higher quality streams (e.g. 4k, 8k), and future-oriented products that may have more in common with video games than traditional broadcasting. The Agile Media Blueprint provides a way forward for media companies.

As the AMWA discusses the points that will be presented in this discussion paper, additional areas that we anticipate discussing include, but are not limited to:

- Scalability and performance of cloud platforms
- Operational validation
- Use of big data and data essence
- Extensibility
- Tooling
- Cloud vendor/on premises portability
- Dealing with continually changing technology and requirements

Architectural Overview

The Agile Media Blueprint is based upon an architecture that captures, manipulates and transports essence streams in the form of data “grains”; small units of video, audio, and data essence such as teletext or closed captioning. Cameras, microphones and other sources emit flows of grains onto a network. Each grain is stamped with a Precision Time Protocol (PTP) timestamp, recording the time at which the visual image or sound wave was captured. Each grain is also given a unique and persistent identity which is permanently tied to, and associated with that grain. Monitors, speakers, and other receivers can reproduce flows of grains as images, sound or data. (Readers may find it helpful to refer to the Agile Media Blueprint diagram in the appendix as they read through the following material.)

Grains are persisted (kept) in an object store. This is the same generic storage technology used by millions of Internet applications today. A very fast, very large cache sits above the object store, temporarily storing and serving out essence that is requested frequently. Essence

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

that is requested less frequently is retrieved from the object store, since it is not found in the cache.

Workflows are accomplished by chaining together the required composable atomic functions. As figure 1 below shows, these functions are small bits of software (atomic) that take in content, perform a transformation, and emit the transformed content. A transform could be a fade between two video sources, it could be a key function that superimposes one video on top of another, or any one of a number of other functions we commonly perform on content. Each function is given an identifier which allows it to be unambiguously referred to and reused. This is shown by the small “ID” tag attached to the function in the figure.

We believe it should be possible to develop an open, collaboratively developed API (a content API) that serves as a consistent interface between the function and other parts of the system. The software interface for all functions would be through this API. This is shown at the bottom of figure 1.

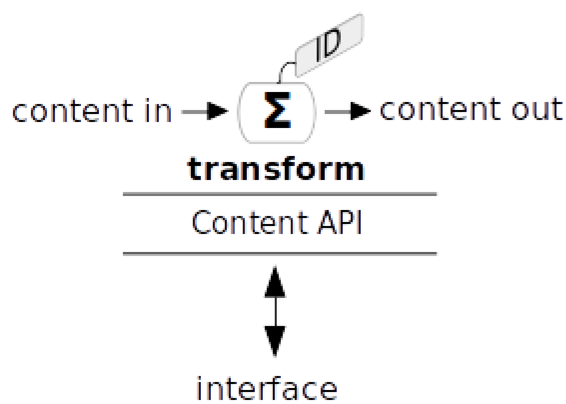

Figure 1. Composable Atomic Functions

Importantly, the function exists by itself. It has no knowledge of other functions, and it makes no assumptions about the workflow either preceding or following it. In this way, the service this function provides stands on its own, and may be used in any number of workflows where the function is needed. The notion of putting together a number of different functions is referred to as composability, and chaining together these different functions is sometimes referred to as service chaining.

Service chaining-
As figure 2 shows below, functions ‘a’, ‘b’, and ‘c’ are chained together, perhaps in a master control switcher. Function ‘a’ is a video cross-fade. Function ‘b’ is a key function that superimposes a moving ‘bug’ on top of the primary video feed. Function ‘c’ inserts a watermark for tracking purposes.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

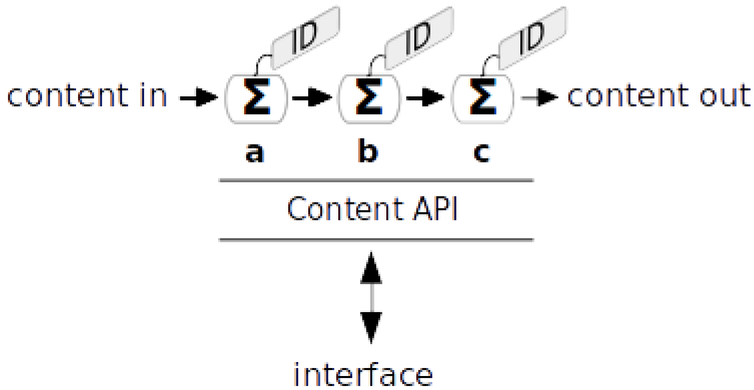

Figure 2. Chaining together a master control switcher workflow.

Taken together, it should be possible to replicate the functionality of a modern master control switcher, for example, as a chain of atomic functions.

When the need arises, it could be possible to chain together another group of functions to perform a different workflow. Some of these functions might be the same as were used in the master control switcher and some might be different. This is shown in figure 3 below. Note that in this case a new function, function ‘f’ has been added. Functions ‘a’ and ‘c’ are reused, but they are in a different order. Function ‘f’ might be a function that synchronizes audio to video. The exact purpose of this workflow is not important. The critical points are that each function performs a small piece of work; that the function stands alone; that it uses a standard interface, the Content API; and that the functions may be chained together to create different workflows to accomplish a wide range of tasks.

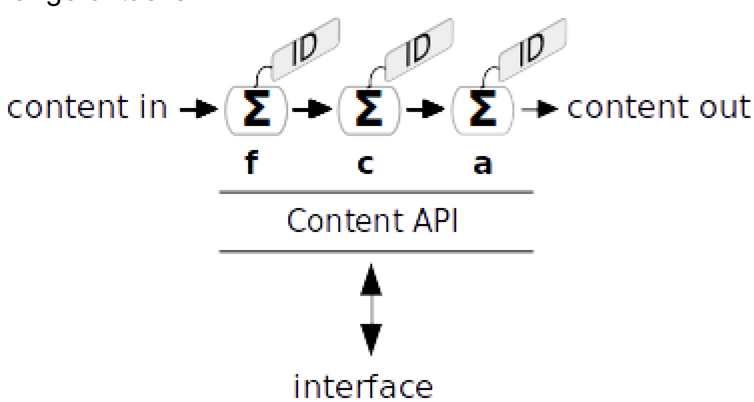

Figure 3. Services may be reused to create different workflows

Service Chain Automation-

One of the strengths of working with composable functions is that workflows may be built and torn down as required. As you might imagine, the construction and deployment of these service chains may be orchestrated using widely-available tools that are used to build workflow chains for the back-ends of major Internet applications. Different approaches to function composition can be applied to enable and optimize media workflows, for example:

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

- Service Function Chaining[1] (SFC) in the network;
- microservices running in containers, such as kubernetes[2];
- software modules linked ad-hoc at runtime by distributed message queues (e.g. reliable queue in Redis[3]).

In addition to building service chains, some automation tools dynamically create new instances of functions in the workflow chain if a function were to stop responding, or if the load on a particular service were to become too high. This dynamic, automated scalability is one of the keys to using professional media in a cloud-fit environment where performance needs to be monitored and actions need to be taken if failures occur.

Enabling technologies

The following are key enabling technologies in the Agile Media Blueprint.

Content API

At the core of the AMB is the Content Application Programming Interface (“Content API”). The content API provides a consistent interface across many different atomic services and storage systems. Using it, applications may read, write and make grains on-the-fly. Any device, from a clip-on microphone to a central sports production facility can expose or request content using this API.

Applications could access the Content API through a URL. An example of the possible formatting of the URL follows. This example is offered as a starting point for discussions within the appropriate AMWA committee:

https://<device>.<domain>/<content>/<stream>/<time>.<format>

The path parts are:

- <device> - a local name for the device, possibly an alias
- <domain> - a global or .local domain name
- <content> - sub-resources that represent complete items of content, with optional extra services that help find the content, such as "search.html" or "content.json"
- <stream> - sub-resource that represents component parts of the content item, such as streams or graphic files, and additional resources that can be used to find, synchronize and process the components, such as "cable.json" or "description.html"

[1] See the IETF RFC at https://tools.ietf.org/html/rfc7665
[2] See https://kubernetes.io/docs/concepts/overview/what-is-kubernetes/ for additional information.
[3] See https://redis.io/commands/rpoplpush#pattern-reliable-queue and search for related packages.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

- <time> - for streams, access to grains or grain ranges identified by PTP timestamps, with relative PTP timings used as the ubiquitous timing model of the API. Other resources at this level are used to describe the stream, e.g. "technical_data.json"
- <format> - a means of describing the format that is being requested, which may include size, scale, codec, wrapping etc.. The format may be described further through headers, query parameters or in the body of a request.

The power of names

Very early on, IP networks had to be manually configured. Each computer was assigned an IP address, and the name, or identifier "Mail Server" was just what everyone called the computer with the IP address 192.168.4.3. But what if the server with that IP address failed? How could you quickly change to another mail server with a different IP address without having to manually configure every computer on the network? The Domain Name System or DNS provided a solution to this problem.

DNS

Names on the Internet are fundamental. These names usually take the form of Uniform Resource Identifiers or URIs[4]. (URIs are a super-set of Uniform Resource Locators or URLs such as https://www.amwa.tv.) The name (sometimes called an immutable identifier) remains the same throughout the lifetime of an object, even as that object changes location and takes new forms. An example of a name could be 'mailserver'.

With DNS, the name 'mailserver' becomes decoupled from the IP address 192.168.4.3. If the mail server with that IP address does down, DNS can automatically redirect systems looking for 'mailserver' to another mail server with a different IP address. DNS may also be used to do load sharing by directing queries to a group of servers using various load-balancing algorithms.

DNS, generally speaking, is an underused resource in the professional media environment. Making use of DNS allows us to employ any number of techniques and best practices for business continuation. Furthermore, the DNS system is continuously monitored and maintained by some of the world's best security experts, and employs cutting-edge security technologies on a global basis.

URLs

Unique Resource Locators (URLs) add a hierarchical path to a DNS identifier allowing us to assign an immutable identity to a resource. In combination with the HTTP, resources can be

[4] See IETF RFC 3986 ***Uniform Resource Identifier (URI): Generic Syntax*** authored by T. Berners-Lee et.al. for more information on URIs.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

organised, published, referenced, downloaded and uploaded. Resources are not just mail servers. A resource could represent:

- people, teams, organisations
- devices - cameras, microphones, speakers, monitors, control surfaces - both with an actual name and an alias for their current use, e.g. “Camera 17” aliased to “Studio A Camera 3”
- streams - content flowing from a source to destination, whether it is being pulled or pushed
- content - collections of streams and stored content with associated data
- compositions - recipes for assembling new items of content from existing ones, just-in-time or just-in-case
- every grain - frame, audio sample, and event

Virtualization

Virtualization is the process of mapping several logical computer systems to one or more physical computer systems. An application can be installed onto a virtual machine, along with just enough of operating system, in its own isolated environment. Multiple copies of the same virtual machine may be deployed on demand, allowing optimization of physical computing resources and providing on-demand scaling.

Caching

In many of the use cases and diagrams that follow, a cache is shown between the Content API and the object store. The cache plays a critical role in making the platform usable for media applications. The cache’s job is to ensure that content is available within acceptable performance tolerances. The cache is also responsible for providing scaling when needed.

A massive amount of investment has gone into developing caching layers and technology for the Internet, with an ever growing number of features. For popular websites, large distributed RAM caches - up to 9Tb in size - become the business logic hub for serving requests, mediating microservices and applying artificial intelligence.

Security

Security is built into the platform, and by extension, the AMB, from the outset. The AMB starts with people and their roles, and leverages best practices in security, using approaches that rely upon national security frameworks and that comply with the latest requirements of corporate and legal security auditing. Security is discussed in greater detail in the section titled, “Introducing the Agile Media Blueprint” below.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

Metering

Metering, which tracks the use of a particular application or function, is perhaps, an unfamiliar concept for some in the media industry, but it can enable new business models for media companies and suppliers alike. New consumption models are also enabled where a media company may choose to pay for the types and amounts of functions they use. This allows them to try out new business ideas without incurring the fixed costs normally incurred in traditional broadcast architectures.

Other key concepts

Other key concepts of the Agile Media Blueprint include:

- Focused on people, content and monetization
- Capable of consuming and delivering SDI, SMPTE ST 2110 files and OTT-style streams
- Fast try/Fast fail - try new ideas quickly and at minimal cost. If they don't succeed, try something else!
- Format agnostic
- Frame rate agnostic
- Geographic diversity & business continuation built-in
- Takes advantage of the large pool of developers and tools for Internet technology

While the Agile Media Blueprint has yet to be demonstrated at scale, it has the potential to offer significant benefits to broadcasters at a time when they are faced with serious operational challenges. The architecture employed by the Agile Media Blueprint has been proven in mass-market social media applications conducts thousands of transactions per millisecond.

The key enabling technology behind the AMB is actually the technology stack and architectural patterns behind social media, big data, and over-the-top (OTT) video streaming. Massively parallel with fast, non-blocking networking, the platform of the Internet has reached the point where it is now time to begin looking into whether the ideas presented in this discussion paper can be proven out through a program of serious and in-depth testing in a number of real-world scenarios. We propose that the AMWA begin a business use-case driven, future-looking activity to do just that.

The AMB and its relationship with the JT-NM Roadmap

This blueprint embodies the concepts in the Joint Taskforce on Networked Media (JT-NM) Roadmap of Networked Media Open Interoperability (JT-NM Roadmap[5]) specifically with regard to building dematerialized facilities[6]. The roadmap is shown in figure 4 below.

[5] The roadmap is available at https://www.jt-nm.org/roadmap
[6] See http://www.jt-nm.org/documents/DematerialisationBackground_2017_09_05.pdf for an in-depth discussion of dematerialization

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

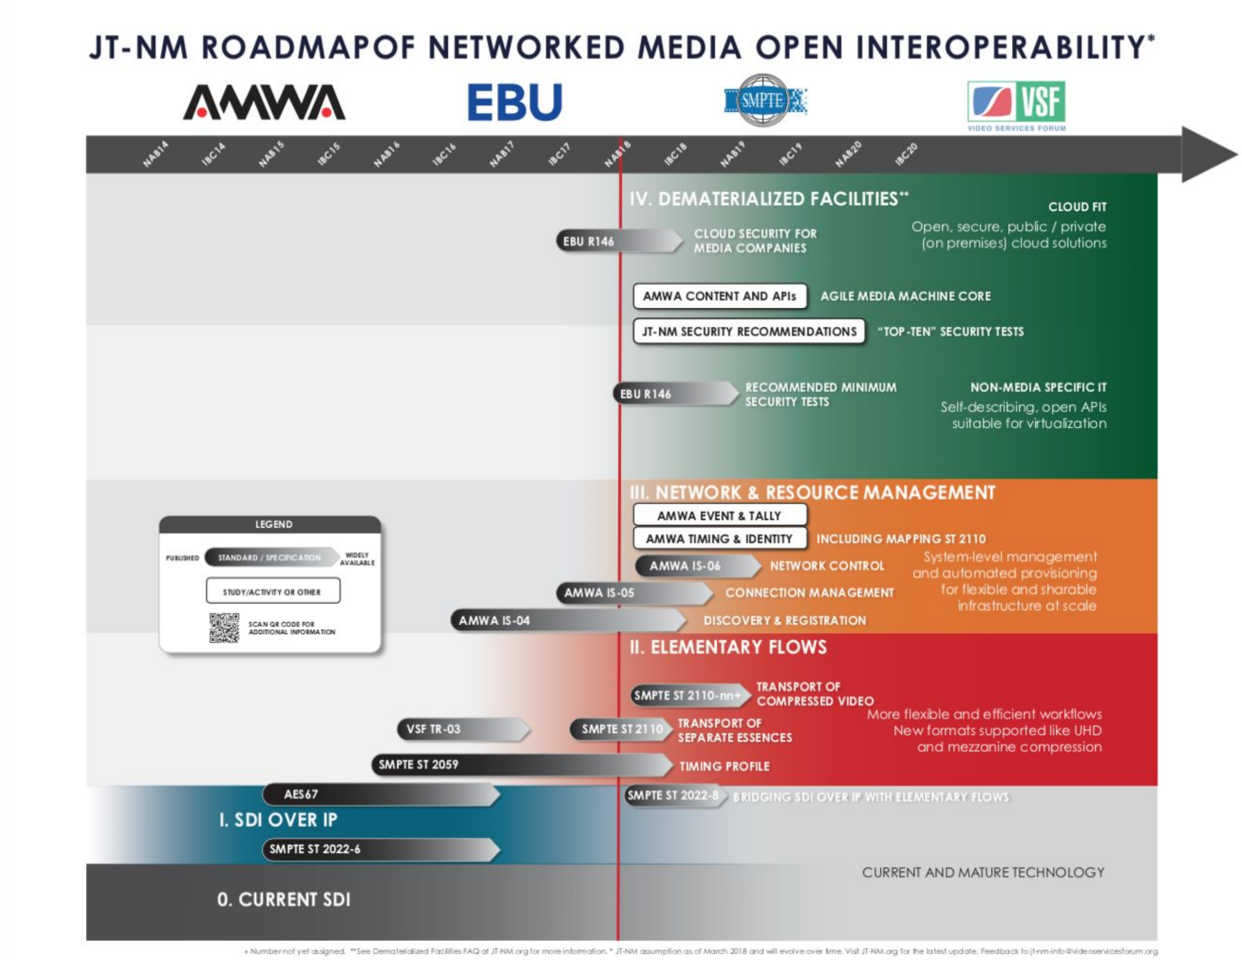

Figure 4. The Joint Task Force on Networked Media Roadmap of Networked Media Open Interoperability

The Agile Media Blueprint and the concepts presented here provide potential direction and concrete suggestions for how the industry might approach developing collaborative standards and specifications specifically targeted at the 'dematerialized facilities' space shown in green in the roadmap.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

Use cases

The following use cases are intended to lay out real-world scenarios and to describe how their requirements may be met by the Agile Media Blueprint. Frequently, when a new architecture is introduced, the first question a reader may ask is, “can I use this new approach to do what I am currently doing today?”. Use case 1 has been developed to address this question.

Use Case 1 - Master Control Room / Live Sport / Live Studio

Current Operation

This use case begins by describing the current operation of a media company with a linear channel that is distributed nation-wide. Most of the time the channel transmits scheduled, pre-produced programs combined with short promotions and commercials. The company also deploys OB vans (sometimes called Remote Trucks) to sports arenas where they are used to create a live feed. This feed is sent back to a Master Control Room (MCR) for inclusion in the national broadcast. During the live broadcast, the MCR facility takes additional live feeds from commentators on a stage located in a local production studio. This local studio programming is broadcast before and after the remote live event. When the live event is finished, MCR again plays out pre-produced programs. A simplified diagram of this use case, using conventional SDI and MXF technology, is shown in figure 5 below.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

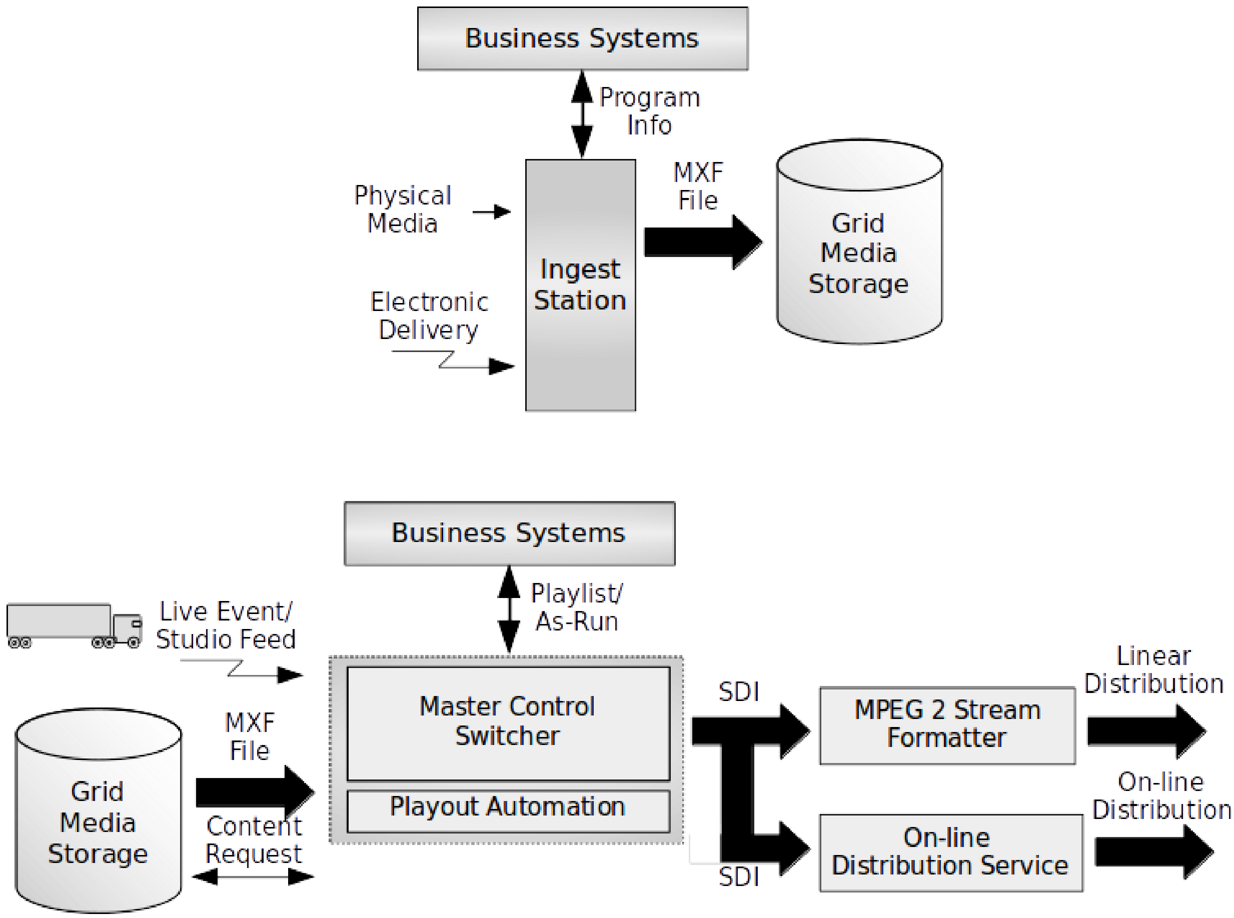

Figure 5. Traditional SDI/File-based Master Control Facility

Starting at the upper-left, pre-produced content is delivered to the facility either on physical media or through a file delivery service. Physical media is brought to an “ingest station” where a computer workstation with associated software is used to transfer the contents to an MXF file, assign an ID, and to populate a database with necessary metadata. The ingest system interfaces to business systems, exchanging information about the content that is expected for delivery, assigned identifiers, and the status of the content. The MXF file created by the ingest station is stored on a central grid media storage system in a proprietary internal format.

The MCR area, shown in lower portion of the figure, consists of the same grid media storage system, a master control switcher and a playout automation system. The automation system interfaces to business systems and receives a schedule of material for playout, arranged according to time; called a playlist. It also exchanges real-time updates with the business system throughout the day. The automation system interacts with the server and the MCR switcher at air time causing the appropriate content to be played, and causing the switcher to take that content and emit it on the linear output.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

The output of the switcher is an SDI signal that is converted to an MPEG-2 Transport Stream. This stream is used to feed various linear distribution channels such as a transmitter or a cable system. This same feed is sent to an outside on-line media company for Adaptive Bit Rate (ABR) “chunking” and distribution via the Internet.

From time to time, the media company has the rights to broadcast live sporting events. Before and after the remote live event, the MCR facility takes a live feed from a local studio facility. During the live event, the MCR takes a feed from the OB van. MCR inserts commercials and promotional material using “live assisted automation”.

Future Operation

Let us look at how this same use case can be supported in a future facility implementing the Agile Media Blueprint. Again, the point of this use case is to answer the question, “can I do what I do today using the AMB?”.

The upper left corner of figure 6 below is essentially identical to the current facility, with physical media and electronically delivered programming going into an ingest workstation. The ingest station still exchanges information with business systems. However, there are two fundamental differences in this workflow. First, as content is ingested, it is converted to grains, and the associated identity and timing metadata is permanently bound to the content. Second, the content is stored as grains in a generic object store rather than in a proprietary file format on a specialist video server. Importantly, all interaction with the object store is through an open and publicly available Content API.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

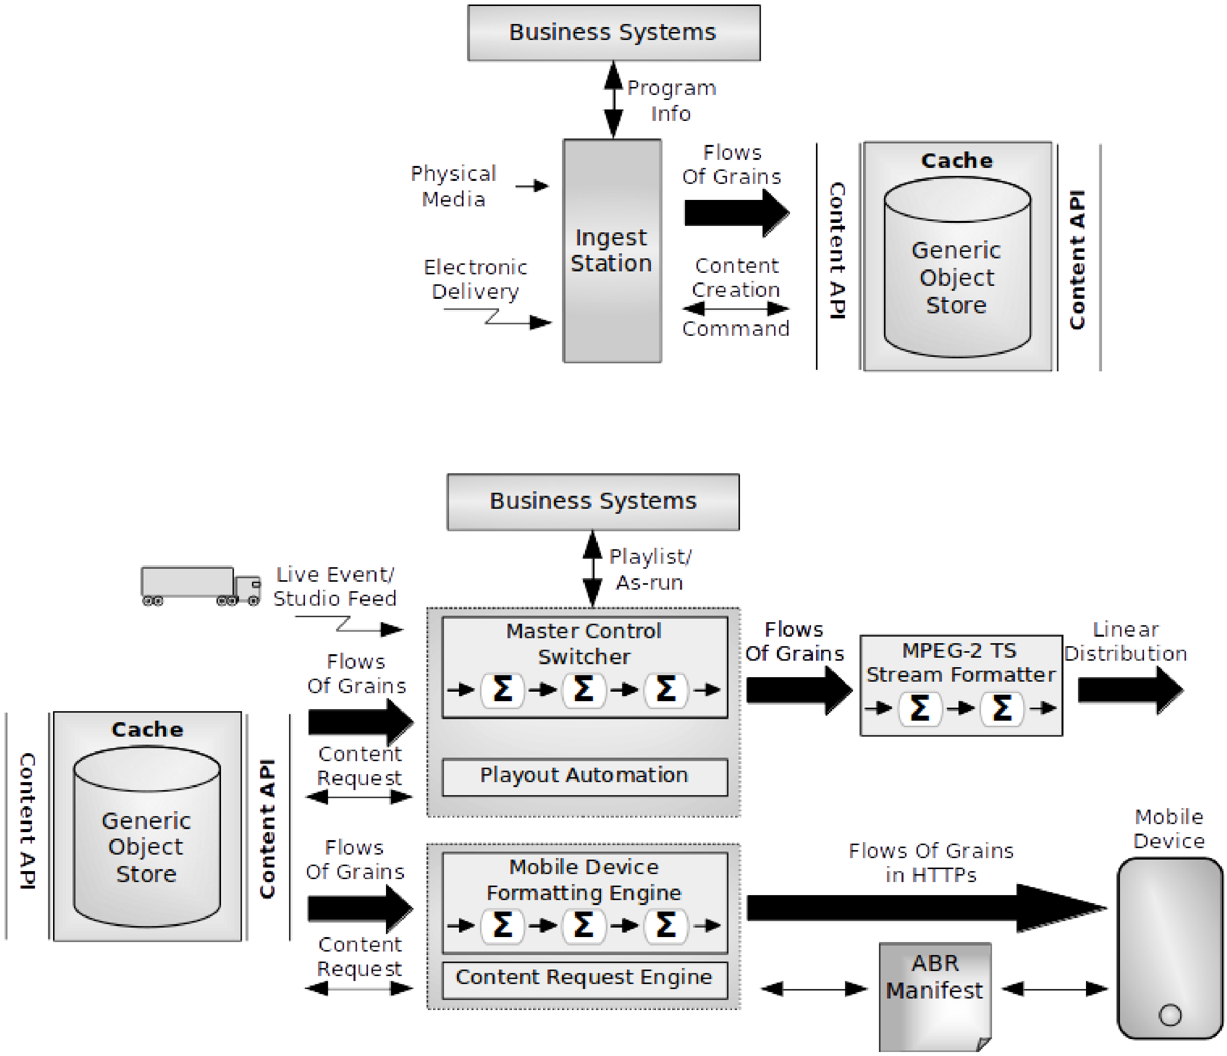

Figure 6. Agile Media Blueprint Master Control Facility

As in the SDI case above, the MCR facility is shown within a dotted box in the center of figure 6 above. The MCR switcher consists of a number of functions chained together to achieve the same output as that obtained by the traditional MCR switcher. The automation software interacts with the object store through the same open Content API that is used during the ingest step. The switcher emits flows of grains, which are sent to an MPEG-2 TS stream formatter, which also consists of a number of atomic processing functions. From there, the stream is sent to a transmitter or to other traditional linear distribution points.

As in the current case above, It would be possible for the linear feed from this AMB facility to be sent to an on-line service to be chunked and distributed via the Internet. But the AMB can provide an attractive alternative. Due to the similarity of the atomic processing functions of the AMB and adaptive bit rate transcoders, it is possible to efficiently distribute the Internet feed from the same facility used for the linear feed. This provides a cost saving by bringing distribution stream processing in house. As an additional benefit, using the same facility, the

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

media company can have a more personalized relationship with its audience. It becomes possible to stream personalized content to mobile devices and smart TVs using content that is either stored in the object store, or being streamed live. Other scenarios are possible with this infrastructure. See Use Case 2 below.

In figure 7, the live studio and all of the broadcaster's OB vans have been converted to Agile Media Blueprint technology. They emit flows of grains that are transported using HTTPS protocol to the object store. As soon as grains are written, they are available. The automation system requests these live grain flows from the cache, once again leveraging the Content API. From this point, those flows move through the workflow established by the MCR switcher and out to distribution points.

Critics may point out that the AMB cannot deliver the low latency required in this use case. But it is worth noting that given current technology and compression, “live” broadcasts as they exist today are far from live. Also consider that HTTPS can provide faster than realtime transport. We concede that there may be cases where the performance of technologies inside the platform may adversely impact real-time workflows. In this case, an AMB facility can sit alongside a traditional broadcast plant, providing enhanced capabilities to the broadcaster, and importantly, providing a migration path from SDI to the future.

This use case started with a traditional SDI workflow, and has shown how the same workflow may be supported using an Agile Media Blueprint design. While the move to IT changes much of the underlying technology base, those changes are abstracted away from the broadcaster, allowing us to present this use case in a way that is very familiar.

One may question why go through the changes necessary to implement the AMB when there is almost no change in workflow. The answer is that once the facility has incorporated an AMB architecture, many new use cases can be addressed. Please see the use case below for examples.

Use Case 2 - Live Broadcast With Integrated User-Contributed Content

This use case starts with the facility described in use case 1 above. But in this use case, the media company has also secured rights that create new revenue streams. These opportunities are as follows:

- User-contributed content; live content from user devices is integrated into the stream created by the broadcaster
- User-contributed live event sports production; previously vetted and authorized end-users are allowed to “customize” their own live event broadcast based upon content assets made available by the media company. The public may choose to “follow” this end-user, and may also be allowed to “join” these broadcasts.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

- Artificial Intelligence (AI) live event production; an AI system that is aware of a viewer's preferences creates a personalized live event, showing highlights of the viewer's favorite athlete, making available additional interstitial content featuring that athlete or his team and providing commercial opportunities to advertisers during breaks in live action at the venue.
- The OB van(s) are eliminated; live production is done in-house at the media facility.

User-Contributed Content

Referring to figure 7 below, in addition to cameras and microphones provided at the venue by the broadcaster, user-contributed content is available through apps written for mobile devices. Users running these apps, and with appropriate authorization, become potential sources of content to be integrated into the live broadcast. The apps on the devices produce HTTPS flows of grains with the required identity and timing headers.

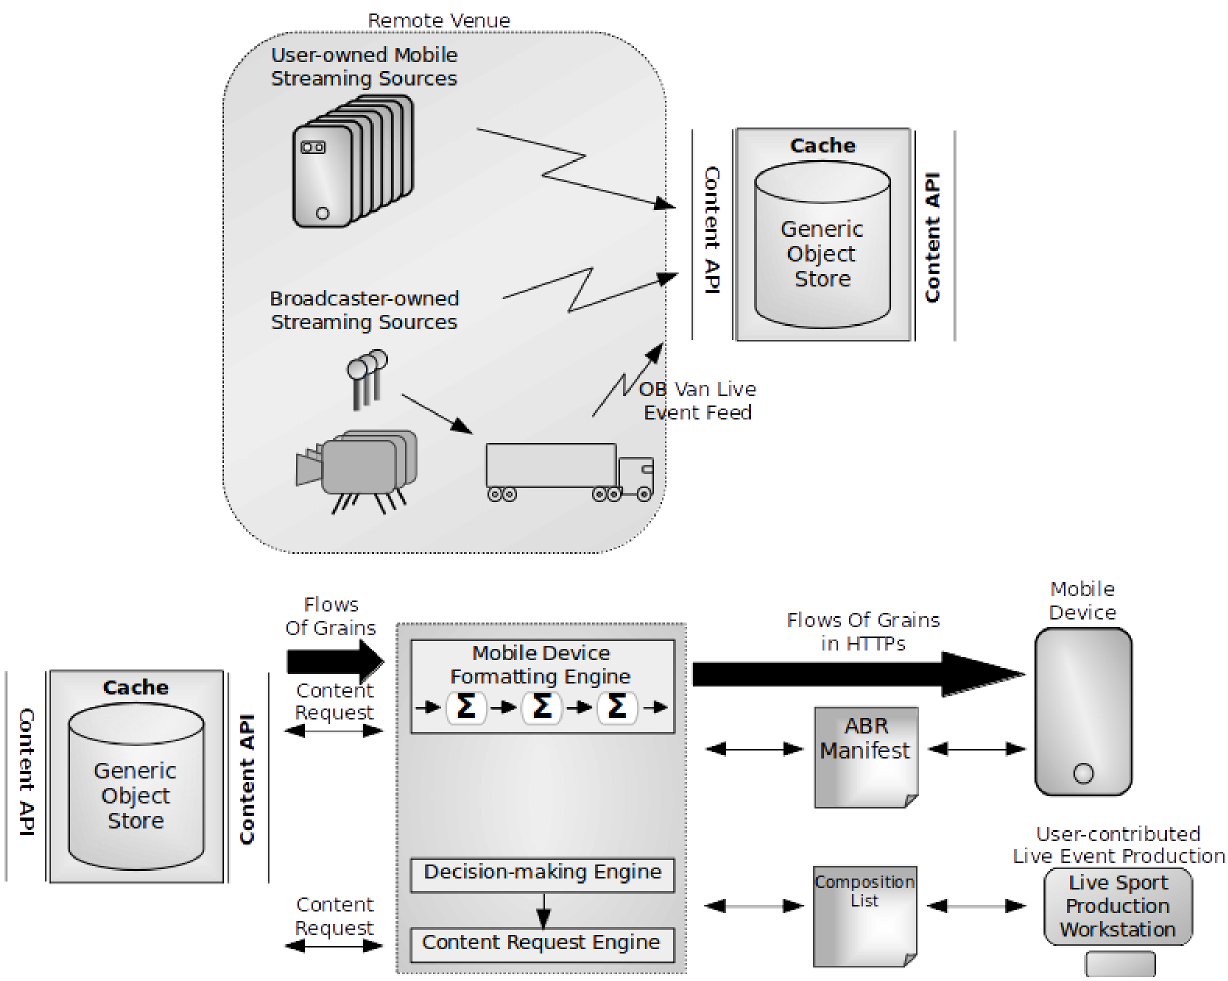

Figure 7. User-Contributed Live-Event Content and Remote Sport Production

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

User-Contributed Live Event Sport Production

As shown in the lower right portion of figure 7 above, a "remote producer" creates a new "live event feed" using a Live Sport Production Workstation. With the proper authorization, that producer interacts with decision-making engines and content request engines to access live streams from the generic object store. The decision-making engine and the content-requesting engine analyze the composition list and then generate content requests that are appropriately authorized and that meet the technical requirements, e.g. resolution, available bandwidth, etc.

The application used by the remote producer interacts with the object store and clustered cache through the Content API, allowing him or her to select different camera sources. There is the opportunity for suppliers to create new software tools that allow the remote producer to add graphics, perform video effects, and to do other tasks necessary in a live production.

The remote producer promotes his or her channel through social media (perhaps with the assistance of the Broadcaster), and viewers may "follow" the producer and join the live events he or she creates. The media company is free to create new business models that potentially include the remote producer, those end-users creating content at the venue, and commercial sponsors.

AI-Created Live Event Content

An AI system running on the platform, having analyzed an end-viewer's preferences, provides a personalized live feed focused on that viewer's favorite athlete. The AI accesses the object store through the Content API, and interacts with the end-viewer. That end viewer has the option to give input to the AI, "liking" and "not liking" portions of the feed. As a result, the AI learns and delivers a feed that reflects the viewer's preferences at that moment. This scenario is shown in figure 8 below.

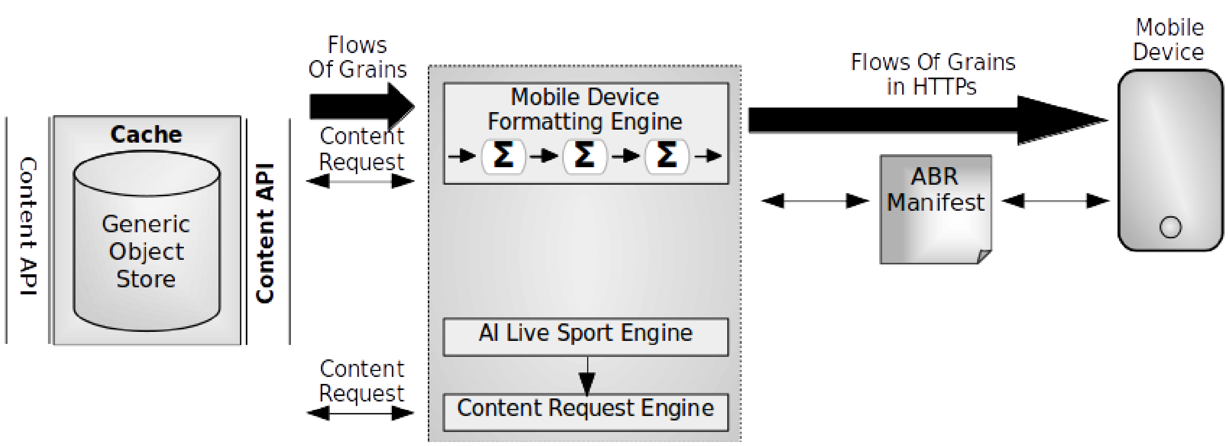

Figure 8. AI-produced Live Sport Engine

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

Local Production of Remote Events

In this use case, given the capabilities described, it becomes feasible to consider eliminating OB vans entirely. High-bandwidth connections between venues and media facilities allow the "backhaul" of all cameras, microphones and other sources, allowing the media company to produce the live broadcast in a fixed facility rather than in a OB truck at the venue itself. The cost savings represented by this change could be significant, as could the increase in the number of live events it is now possible to cover. This backhaul scenario is shown in the upper portion of figure 7 above. Note that this use case is not as unrealistic as it may seem. The Swedish national broadcaster SVT has been pursuing this approach for several years[7].

Use Case Conclusions

These use cases may seem unusual. The fact is that with regard to live sports, entirely different products are rapidly evolving[8] and becoming ever more social-media-like. A broadcaster with a conventional SDI plant would not be able to take advantage of the business opportunities presented in use case 2. While we do not know if these specific use cases will arise in real business operations, applications such as PunditSports and Periscope[9] are redefining the meaning of live event video. The Agile Media Blueprint provides a way forward for media professionals who seek to provide competitive products in this application space.

[7] See article, *"SVT Discusses Workflow Innovation Ahead of the Stockholm Open in SVG Europe."* https://www.svgeurope.org/blog/headlines/svt-discusses-workflow-innovations-ahead-of-the-stockholm-open/ accessed 2 April, 2018
[8] See https://www.punditsports.com/ as an example of user-contributed content and live sport
[9] http://www.adweek.com/digital/periscope-now-allows-android-users-to-broadcast-live-360-video/

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

Introducing The Agile Media Blueprint

The Agile Media Blueprint starts with people. These people may be working in teams, either within a single organization, or across organisations. As shown in figure 9 below, it begins here, not only because people are a key component in our craft-oriented industry, but because by starting here, one can address security at the outset. Administrators can assign roles to people and grant them permissions based upon those roles. People in a different organizations may be given specific rights in your organization based upon business needs. For example, personnel from two remote production companies using two different OB vans may be given roles that allow them to collaborate for a one-time event. This architectural approach allows the system to grant people access to physical resources across a shared content API, based upon their roles and organizations.

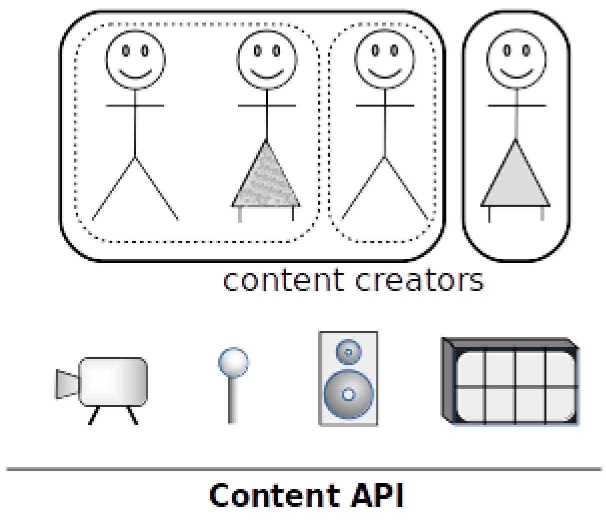

Figure 9. People And Their Roles Are A Key Part Of The Agile Media Blueprint

While the Agile Media Blueprint can be deployed to meet professional media production needs, it also may be used to efficiently produce new types of content that require direct interaction with end viewers. As illustrated in figure 10 below, the Agile Media Blueprint targets social consumers, through broadcast media, OTT and mobile.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

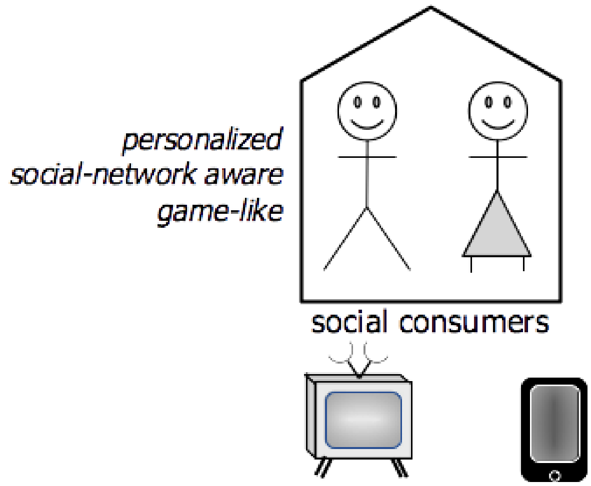

Figure 10. Consumers Interact With Bespoke Versions of Content

Conventional broadcast facilities could never provide individualized content to tens of thousands (millions?) of viewers - it is an anathema to the fundamental concept of “broadcasting”. But to deliver whatever a viewer wants to see when they want to see it means that it is absolutely critical that AMB facilities scale seamlessly. In many cases, in an “individualised content” scenario, content is cached and streamed to the viewer as a one-off event. Many broadcasters are already paying for bandwidth at a cost scaled per viewer, without benefiting from a closer, bidirectional relationship.

As figure 11 below shows, at the core of the Agile Media Blueprint is a grain-based Content API for element-by-element access to the media, backed by best-of-breed cloud-fit technology, including RAM-caches, AI and object stores. Microservices make new media elements to order as they are required; no more wasting time and resources by creating media elements “just in case” they are needed.

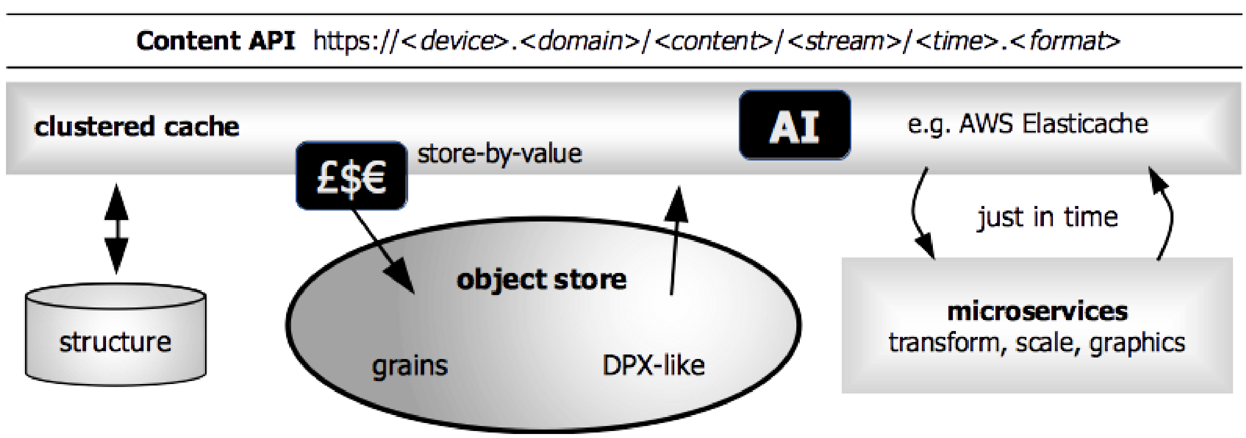

Figure 11. A Common Content API Provides A Uniform, Open Interface To Essence

Media transport in the Agile Media Blueprint is between Content APIs, with bidirectional links

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

operating in parallel, faster than, slower than, or at real time. As shown in figure 12 below, a time-to-bytes translation interface (bytes-to-time component) allows content to be read and written to common signal and file formats, and as described above, allows for compositions to be made on-the-fly.

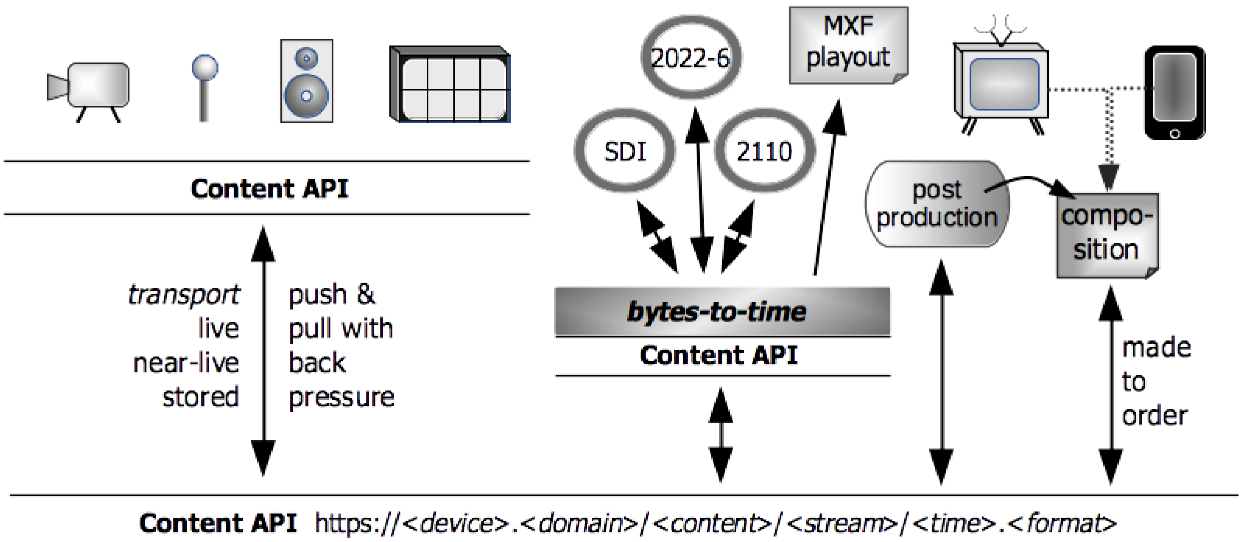

Figure 12. Bytes-to-time Conversion Allows Creation/Consumption Of Existing Formats

Libraries of files and streams can be migrated into an implementation of the Agile Media Blueprint, either just-in-time, or according to a schedule, using a bytes-to-time unwrap component, supported by an index database (see figure 13 below). This allows media factories to easily migrate from a file-based infrastructure to an API-based infrastructure.

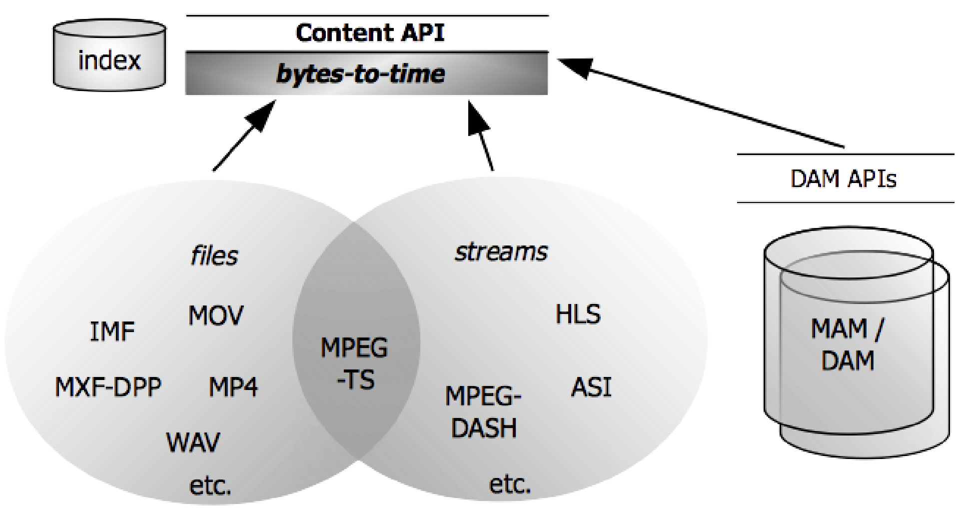

Figure 13. Bytes-to-time conversion and an Index Facilitates The Creation Of Common Media Files and Stream Formats

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

Security

IT already has excellent mechanisms for managing users, groups, organisations, roles etc. These mechanisms support the safe sharing of resources between people and across organisations. IT also provides a set of highly scalable and securable protocols for data transport that have survived the evolution of the Internet. Until now, the media industry has not been able to leverage these mechanisms, largely because media has existed outside of the IT world as opaque serial streams (SDI) or in non-IT-friendly file formats (MXF).

Because our industry evolved separately from the IT ecosystem, security was frequently “bolted on” to our media applications. And because we built media systems with bespoke products from a multitude of vendors, we seldom got a joined-up approach to security. This left us in the unenviable position of attempting to create a hardened, cutting-edge security environment out of a patchwork quilt. This is a difficult, and some would say, impossible task.

To be fair, the early Internet evolved without a lot of attention paid to security, and in a largely collegiate environment (that said, the United States military initiated development of the Internet with very specific security requirements in mind). But two things have vaulted IT technology ahead of bespoke media platforms.

- Malicious people appeared pretty early in the evolution of the Internet. Bad guys, or perhaps just curious college students with a lot of time on their hands, started picking away at any server they could reach. This means that IT departments have been dealing with security threats for a long time, while Broadcasters remained insulated from the fray, at least until relatively recently.
- Because platforms such as Amazon and Azure provide a joined-up approach to security, developers on these platforms benefit from a common framework across a broad range of applications. This is a relatively recent development, springing from the fact that many cloud platforms offer a homogeneous approach to security that is then employed by application developers writing for those platforms. Over time, this change has allowed IT security professionals to develop a set of best practices, technologies, and even security-specific jobs around this topic.

While broadcasters and media companies might wish to “air-gap” themselves from the Internet and the security issues it raises, this is just not practical. Connections, and security liabilities, are inevitable for any company that wishes to actively engage in the media industry in the modern age.

Increasingly, government regulators and investors are requiring that media companies perform security audits. Suppliers may soon find themselves needing to provide certifications stating that their products conform to certain best practices, Standards and laws. This trend is

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

accelerating, especially as more countries identify broadcasting and media as critical national infrastructure.

Finally, it is important to consider security from a developer's point of view. In the end, it is the developer that is charged with finding a way to produce applications efficiently while at the same time meeting increasingly stringent security requirements.

Business Summary

Moving to the Agile Media Blueprint presents a number of challenges to us from a business perspective, including changes to mindset and a fundamental change in the underpinning of the technology basis of our business. But as the discussion above has pointed out, there are major benefits to adopting this approach going forward. That said, two key business-related questions remain. First, how can I make money with this new technological approach, and second, is it real?

Cost effective resource management

Media companies can make money with the AMB by being able to:

- leverage the investments being made in the platform
- rapidly scale up and out, or down and in depending upon demand
- use functions and capabilities on an as-needed basis
- deploy an infrastructure that has security built in from the beginning
- create new products that provide interaction with end viewers, resulting in personalized content delivered to them in ways that could never be done using conventional media platforms
- Make better use of capital investment. For example, in a traditional infrastructure, technical resources that might be needed in New York might be deployed in San Francisco. The AMB provides access to resources independent of the physical location of those resources.

From the outset, the platform has provided monetization opportunities for IT vendors. Suppliers of media applications can take advantage of this fact. Metering provides a pathway to "de-risking" and sharing risk in media partnerships.

Wikipedia defines Serverless computing, in part, as, "a cloud computing execution model in which the cloud provider dynamically managed the allocation of machine resources." "Pricing is based on the actual amount of resources consumed by an application…"[10]. The AMB leverages many characteristics of serverless architecture. For suppliers, this means applications that are

[10] From https://en.wikipedia.org/wiki/Serverless_computing accessed February 7, 2018

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

built to the Agile Media Blueprint are metered, and that new business arrangements charge for capabilities provided by the software, with use-based pricing. Media companies may be incited to look at this pricing model because it allows them to buy what they need and pay for what they use rather than having to pay a certain fixed amount in order to stand up a media-oriented service. This could be especially important where the user has an opportunity that they would like to explore, but they are unsure whether the venture will succeed.

The technology needed to build the AMB exists in some form or another. Many large media organizations are already using cloud infrastructure to make and deliver media at scale (e.g. BBC Media Factory, Discovery). Organizing the components so that the AMB can be delivered as described is work in progress, as is measuring its performance and calculating the relative business value of particular use cases in a revenue-based model. It would be a useful exercise to develop a set of specific use cases and then develop a set of proof points to identify potential challenges.

The purpose of this white paper is to introduce the Agile Media Blueprint and the context within which it exists. As next steps, the Content API presented here will be introduced into the AMWA technical process for discussion and further development/improvement. As mentioned above, it is hoped that projects on object-based media composition could be folded into this work. We are at a point where a deployment of these concepts in a media company is not only possible, but it would accelerate the work presented here significantly.

Readers should note that a number of media companies are already deploying facilities in the cloud, although none that the authors are aware of are using this grain-based approach with an open API. This indicates that the time is right for moving forward with this work in an open and collaborative environment such as the Advanced Media Workflow Association.

Re-use of SMPTE ST 2110 infrastructure

SMPTE 2110 is a good first step towards building IP infrastructure inside a facility. It enables the transition from point-to-point SDI to IP networking and elementary streams. It also supports more efficient carriage of media data payloads without blanking intervals. ST 2110 is a good choice in scenarios where network QoS can be assured, and where carefully managed hardware devices connect to a network.

ST 2110 devices can be used at the edge of fully virtualized AMB environments. They may also be used in applications where network management is outside the control of the user (e.g. in the cloud). But ST 2110 is difficult to read in modern software, and in cloud environments ST 2110 can be expensive or impossible to write. This is because 1) sufficiently accurate clocks may not be available in generic IT facilities and, 2) virtualized systems are non-real-time. However, with appropriate buffers and hardware/software, ST 2110 can be used as a gateway to the AMB.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

Investment in ST 2110 technology is a move in the right direction. ST 2110 is complementary to the AMB Content API, and a step toward building a de-materialized facility.

Conclusion

The AMB exploits a different kind of platform from today's facilities and will require a different way of thinking about how to execute current workflows. Do the large files and monolithic batch processing of current file-based media facilities transfer well to the cloud? Can you or should you replicate a precisely timed and highly synchronous point-to-point signal using a packet-based network?

By representing and manipulating media as grains synchronized by timestamps, you step away from the need to precisely time the flow of a linear bitstream down a cable, and into a structure for media that is much easier to transport, store and process on modern computer systems.

Non-real time, but not slow, these systems are asynchronous and massively parallel. They are connected to a collection of non-blocking networking optimised for thousands of concurrent connections. These networks contain computers with tens of CPUs, to thousands of GPU cores in each computer. This sort of processing power is intriguing - for example, with a grain-based model and enough computing resources, whole feature film could be transcoded in parallel in a few seconds!

Whether you have performance sufficient for live becomes a question of what delay is acceptable and whether the vast majority of the grains are being transported faster than real time[11].

Many file-based workflows are batch-based, with processes that start at the beginning and read to the end. A grain-based approach is radically different; using a bytes-to-time API, file systems and object stores can provide random access into the grains wrapped inside each file. Files can be viewed as elementary streams, and in combination with virtual composition lists and efficient redirection, the versioning benefits of advanced containers such as IMF are achievable.

TCP, with its congestion control algorithms optimised for use within a facility, rightly gained a bad reputation for wide area network linear file transfer on congested networks. This created a market for UDP-based file accelerators. With major advances in TCP technology, TCP is now proving to be adequate for real time and even faster than real time media transport. This offers the option to do content creation and transfer just-in-time rather than file movement just-in-case. Moving files "just-in-case" is not only costly in terms of people and resources, it is also largely unnecessary in an AMB world.

---

[11] See R I Cartwright. *An Internet of Things Architecture for Cloud-Fit Professional Workflow*. Proceedings of the SMPTE Annual Technical Conference, October 2017. (Also SMPTE Motion Imaging Journal, June 2018, to appear.)

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

In summary, many exciting opportunities for suppliers and media companies exist when the Agile Media Blueprint is leveraged to produce new media experiences.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

Appendix - Agile Media Blueprint

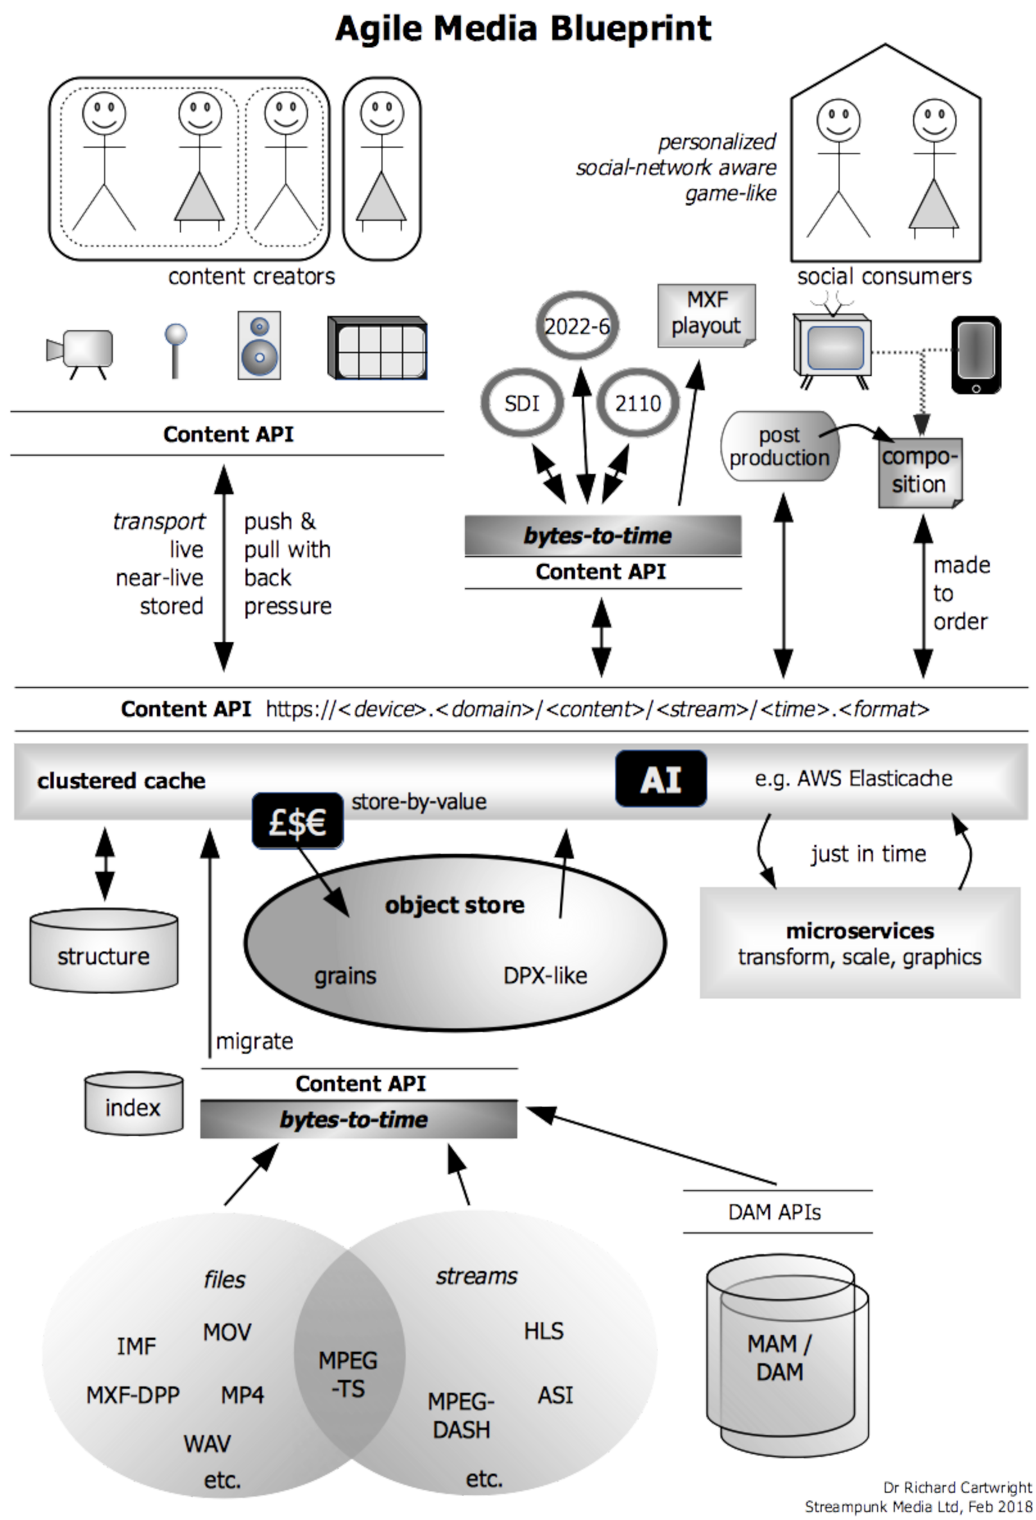

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0

Biographies

Dr. Richard Cartwright - Streampunk Media, Ltd.

Dr Richard Cartwright is CTO and founder of Streampunk Media Ltd. Richard holds a PhD in Computer Science from the University of Warwick, UK, where he studied that application of parallel computation to virtual reality models.

Richard has previously worked at the BBC, Snell Advanced Media, Red Bee Media and was Technical Steering Committee Chair of the Advanced Media Workflow Association. Now Developer Representative on the AMWA board, Richard is a co-author of the Joint Taskforce for Networked Media Reference Architecture and contributor to the JT-NM Roadmap.

He is a prolific developer of open source software for technologies including MXF, AAF, MPEG-TS, SDI and IoT workflows, in languages including C++, Java and Javascript. His ambition is to democratize professional media production ready to deliver a new breed of immersive social television.

Richard is an amateur musician, narrowboat captain and lumberjack at a community-owned forest.

Brad Gilmer - Gilmer & Associates, Inc.

Brad Gilmer is President of Gilmer & Associates, Inc. He is a founding member of the Joint Task Force on Networked Media, Executive Director of the Video Services Forum (VSF), and Executive Director of the Advanced Media Workflow Association. Brad is also the Executive Director of the IP Showcase, an event held at both NAB and IBC, introducing and explaining the IP transition to the industry.

Brad is a SMPTE Fellow and the first recipient of the SMPTE Workflow Systems Medal.

Brad was previously employed at Turner Broadcasting System in Atlanta where he and his staff were responsible for Engineering and Operations for the Entertainment Division Worldwide.

He is an author and editor, contributing two times to the NAB Engineering Handbook, and served as Editor-in-Chief of the File Engineering Handbook. He has written many articles on computers and networking for industry publications.

Brad as an expert witness in media-related court cases at the national level in the United States.

Finally, he is a national climbing and rappelling instructor, pilot, and blues harmonica player.

© 2018 Advanced Media Workflow Association
Creative Commons CC BY-SA 4.0