SMPTE

MOTION IMAGING JOURNAL

Covering Emerging Technologies for the Global Media Community

IP Media Infrastructure—Opportunities, Observations, and Challenges

By Gordon Castle

Introdução:

O artigo desta edição trás as emoções, duvidas e desafios enfrentados por quem pôs a mão na massa para transformar seu fluxo de SDI em IP, então se prepare para ler sobre os desafios de custos associados a adoção do padrão ST 2110, a necessidade de desenvolver sistemas de gerenciamento de mídia como se fossem sistemas de TI. Toda uma transformação tecnológica e operacional enfrentada pela Eurosport (Rede de Televisão esportiva pan-europeia, com sede em Paris) durante sua transição de SDI para IP. Então bora ver o sofrimento alheio? (e aprender com ele, é claro!)

Tom Jones Moreira

As guest editor, I specifically reached out to Gordon Castle for a contribution to this issue of the journal because it is important to hear from different voices as we work our way through the transition from the serial digital interface (SDI) to Internet Protocol (IP) systems. Rarely, if ever do we get to hear directly from the business/technology people who are making key decisions as we transition to IP. He has recent, direct relevant experience as one of the key decision-makers for a major refresh at a large live sports facility in Europe. He is a senior technical strategist with years of experience working both on the user side at Cable News Network (CNN) and later Eurosport and also on the vendor side at Ericsson. The views and opinions in this article are based on his years of experience, and the points he raises are ones that we, as media engineers and experts, should ponder, both with our employers and in society. Castle asks a number of questions in this article. He answers some of them but leaves others as open challenges. Perhaps, someone would like to tackle one or more of these in an upcoming issue of the journal?—Brad Gilmer, Guest Editor

Abstract

While head of technology and operations at Discovery Europe Middle East and Africa (EMEA), I led the planning and implementation of what I believe is the largest all SMPTE ST 2110 Internet Protocol (IP)-based platform in the world. The scale, flexibility, and stability of the infrastructure are terrific, but making it a reality was hard. In this article, I will provide some insights on the goals and the challenges and then raise some questions for the industry as we think about moving live sports production media workflows to the cloud.

Keywords

Digital transformation, Internet Protocol (IP), media infrastructure, ST 2110

Introduction

The story starts in 2018 with the modest goal of upgrading Eurosport's Paris production hub. Like many media production facilities, the infrastructure was serial digital interface (SDI), very purpose-built, and working well. However, it was not meeting the needs of the business as Eurosport looked to provide increased localization, deliver exponentially more content for digital, and improve efficiency. This is, of course, a common problem in media companies every 7–10 years, when a major upgrade is needed. Throughout my career, I have used major upgrades as an opportunity to change the technology infrastructure to provide greater flexibility and reduce operational costs.

The vision was to allow the flexibility and scale to build virtual facilities on the fly that are "right-sized" for the task at hand. The vision is that it will not matter where the equipment is located; in a Eurosport-owned data center, in the cloud, or in a remote physical facility. Virtual facilities are created as required and made available to skilled staff when and where they are needed.

Requirements, Goals, and Results

As we dove into the requirements and started planning, it became very clear that we needed to build a new infrastructure to support all parts of the Eurosport production, including the ten remotely located market production centers and the London operational center. We gained approval for an investment that created a unified digital and linear production ecosystem and a common content hub to enable the growth of our digital and linear businesses. It is driven by three fundamental business goals.

- Support the evolution of live production with a flexible remote production approach.
- Enable digital and linear growth by providing unlimited access to content for all parts of the business.
- Deliver efficiency and improve the bottom line.

On the technology side, in addition to upgrading the infrastructure and meeting the production requirements, we defined three additional goals.

- The ability to support dramatically more local language commentary for growing digital audiences.
- Location flexibility for operational teams with the data center in a remote location (there was no additional technical space available in the building).
- An infrastructure and a technology approach that would allow a graceful transition of live production to the cloud.

The result is a private cloud for live production. The design approach aimed to centralize uncompressed video flow, workloads, content storage, and technology expertise while providing greater flexibility to the production and operational teams. The vision was to allow the flexibility and scale to build virtual facilities on the fly that are "right-sized" for the task at hand. The vision is that it will not matter where the equipment is located; in a Eurosport-owned data center, in the cloud, or in a remote physical facility. Virtual facilities are created as required and made available to skilled staff when and where they are needed.

Figure 1 shows how Eurosport is looking into the future.

Choice of SMPTE 2110

For the video and audio flows, SMPTE ST 2110 seemed like an obvious choice. We naively considered it leading but not bleeding edge. Public and private plug-fests showed that the standard would allow essence interoperability, and several Internet Protocol (IP) facilities existed at this time, although most, if not all, were not full ST 2110. At the NAB 2018 Show, the chief executive officer (CEO) of a leading equipment provider who was not yet supporting ST 2110 asked me if we were moving to ST 2110 because it was "cool" or because we felt we really needed ST 2110. At the time, I felt very strongly that we needed ST 2110. There were two core drivers behind our decision. Technically, we needed the flexibility to handle audio and video separately so that we could support the number of languages and flexibility requirements. From a business aspect, we felt we needed to align with the direction of the industry. Looking back, I am happy to say that both reasons for selecting ST 2110 have proved to be correct—and necessary.

However, while ST 2110 was needed for Eurosport and has achieved wide-scale support across the industry, our experience identifies several important challenges. First, implementation today is still very hard and depending on the scale—possibly too expensive. While plug-fests have helped vendors with interoperability testing and have increased end-use confidence, they primarily focused on the essence level. Audio and video were flowing from device to device—but there was no testing of defined operational patterns, and vendors have made a large number of independent decisions on the parameters. Our system integrator ran all vendors through a standardized ST 2110 test as part of the project. No vendor passed this test the first time. Each vendor had to make software, and in some cases, firmware changes before they could meet our requirements, which were completely aligned with SMPTE standards.

First Question

So, my first question to the industry is—How can we develop and gain alignment around common IP operating models? In our project, we connected with other media companies such as the Canadian Broadcasting Corporation (CBC), and worked with several knowledgeable consultants in an effort to learn and leverage their previous experience. But there was very little to leverage and while ST 2110 and the work of JT-NM TR-1001 provide structure, they lack the details that are needed to make equipment and software truly interoperate. There are many aspects that are left up to interpretation, and no vendor supports all of the parameters. While this is to be expected of a standard, the lack of common practice or common operational patterns makes implementation much more challenging.

Creating a common operational pattern or model is no easy ask. However, now that several media companies have implemented ST 2110, there is a basis of knowledge that can be leveraged. *Like operational pattern 1a (OP1a) as specified in SMPTE Standards for MXF, we need to use some proven implementation that builds on the good work of the industry groups and industry pioneers. I do*

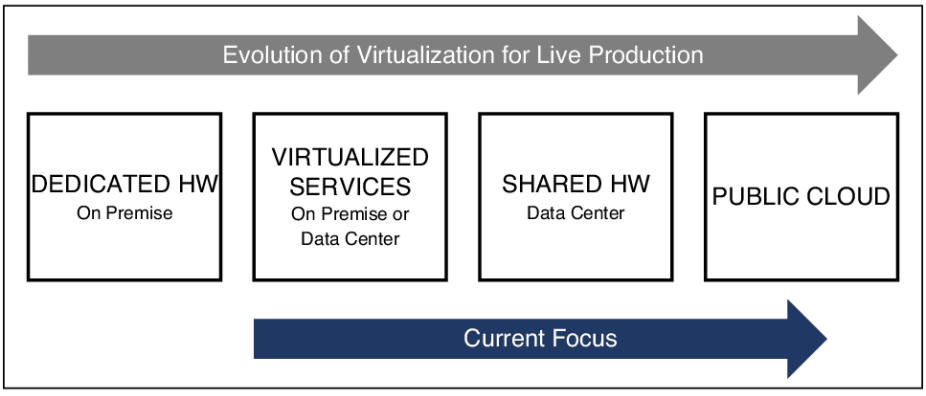

FIGURE 1. Evolution of virtualization and Eurosport's focus.

not know the best way to fund this work but it feels essential to increase adoption and lower the build costs.

The lack of standard operating models leads to the next key challenges—complexity and knowledge. One of the promises of an all-IP base infrastructure is that it is easier to design and manage. You do not have to design all flows in a computer-aided design (CAD) system and build everything around a router that has a size limit. The Eurosport infrastructure has 130,000 endpoints and allows for complete flexibility. Different frame rates, aspect ratios, video, and audio all flow across the same network. But because of the network's size, we did need to use a software-defined network (SDN) and orchestration to manage the network, and we needed dynamic system management tools. But best practice in the industry is still to use CAD for design and to design a fixed implementation that assumes things do not change very often. Just like my home network that started off with a couple of devices and now has over 50 IP devices, networks are in constant change.

Second Question

My second question for the industry is—How can we leverage information technology (IT) tools to manage our broadcast equipment? We started off using Excel which is, unfortunately, still the best practice for managing equipment configuration in a media facility. While this works to get things going—it makes system management and change very complex. Eurosport is now actively moving to an IT toolset that will start replacing the system configuration spreadsheets. This shows great promise but is something we needed to develop ourselves. While I do not think there is a one size fits all, I believe that a more common industry direction will reduce design and implementation costs.

As with the operational patterns, we should leverage the work done by earlier implementers. Both the CBC and Eurosport are using a similar approach although the details of the technology are different. But they are both using IT system tools such as Ansible and Maven to manage media (broadcast) equipment. Looking back, I wonder if the industry needed to develop specific registries such as NMOS IS04 or if a dynamic naming system (DNS) could have worked. *The bottom line here is that these media systems should be designed and managed like IT systems.*

Cost as a Challenge

The overall cost is the next major challenge. There seems to be a common belief that ST 2110 is too expensive for smaller facilities. Given the complexities noted in this article, I agree. However, there is a catch. Industry adoption of 2110 has been very good, which means that it is the only option in some cases. Ironically, there are stories of people edge wrapping ST 2110 into an SDI environment. This is clearly not the goal of an IP transition. But realistically, we must recognize that the IP backplane comes at a different cost point with higher support costs than an SDI router. But it provides great flexibility, allows for A/B network resiliencies, and removes single points of failure (as long as you have multiple switches). Operational patterns, more proven implementation, and more industry expertise will bring costs down and lower the entry points for smaller facilities.

Media companies also need to take a longer-term view on their infrastructure investment. *IP infrastructure will allow ongoing evolution and is not tied to a video format, so it can evolve with the business.* There will still be some major upgrades in our future—but the evolution can be smoother with less dramatic replacement cycles.

NDI and SMPTE 2110

Some facilities are taking advantage of the network device interface (NDI) developed by NewTek, which is a proprietary, but open format with strong vendor support in active use today. It requires NDI connection software that addresses part of the operational patterns issue raised earlier. NDI seems like a great fit for facilities that have limited audio needs or less need for overall flexibility. NDI also provides a much nicer bridge to the cloud as it does not require the timing precision of ST 2110, and the flow data rate is smaller due to compression. There is no attempt in this article to debate ST 2110 versus NDI. Can NDI and ST 2110 work together? In my opinion, ST 2110 provides much greater flexibility for the environment with lots of audio, different formats, and flexibility needs. NDI, however, is less complex and natively works in the cloud. So, along with the operating models—Can there be case studies that show the cost effective use of IP-based systems for small facilities?

Cloud

The cloud aspect of our IP journey is perhaps the most interesting. Many media companies, like Discovery, are running hundreds of channels in the cloud today. The cloud provides amazing scale and a platform that enables ongoing cost reduction and performance improvements. The cloud works well today for content supply chains and channel origination that are focused on producing any number of individual channels. But the needs for live news or live sports production are different. In these scenarios, there is a large volume of incoming material that can be used by control rooms in different markets to produce channels and streams with very low latency. How do we move our live IP media systems to the cloud?

IP is enabling very scalable realtime infrastructure. This is a huge benefit for Outside Broadcast (OB) trucks and live production facilities. At Eurosport, we coupled the IP infrastructure with a cloud stack approach. The implementation is a private cloud approach because there are two data centers that are either owned or rented by Eurosport. But all operational control and production rooms are remotely connected over the Discovery WAN. *This provides centralized content and technology, but distributed operations* (**Fig. 2**).

The private cloud approach also includes a defined technology stack. Media applications sit on top of a platform layer that provides access and provisioning. The platform layer is on top of the infrastructure layer. The infrastructure layer is based on standard IT approaches and includes network management and orchestration. As you work your way up the stack you get to the more media-centric application and in some cases hardware boxes such as video and audio mixers. There are also public cloud applications in the technology stack. As applications move from hardware to software, to software in the cloud, the user impact and transition times can be reduced because they will already be supported by the technology stack.

Architecture

To achieve the flexibility of the vision, it can be useful to think about media applications riding on top of a platform that offers a set of back-end technical capabilities. The applications use these capabilities to perform their tasks. The job of the platform is to deliver the capabilities when and where they are needed, according to a previously established facility schedule (**Fig. 3**).

The platform is comprised of a mix of IP-based broadcast equipment, IT services and applications, on-premises generic IT equipment and networks, and off-premises cloud infrastructure operated by others. This equipment, combined with software in the platform layer, creates the collection of capabilities referred to earlier.

The platform is quite large and is capable of providing a large number of capabilities to the applications above it. The traditional way to build a Eurosport facility would be to combine the application, platform, and physical infrastructures in the physical workspace. However, this creates a vertically integrated and fixed facility. Furthermore, because of the nature of the media business, it results in relatively low utilization of resources.

The new architecture approach is much more dynamic (**Fig. 4**). Resources are scheduled and consumed, and when the event is complete, those resources are released back into the platform to be reused somewhere else. A process of planning, provisioning, and construction enabled by the high-level architecture results in the ability to create media facilities with the correct collection of hardware and software to accomplish the creative task at hand. The architecture presented here employs a set of layers that are intended to provide a clear separation of concerns at the different levels in the system. The layers are planning, provisioning, construction, and control.

Private Cloud

The Eurosport private cloud is constructed from a mix of virtualizable applications (such as edit, graphics control, and facility control) and nonvirtualizable applications (such as vision mixing, ingest, audio mixing, and so on).

Located within the private cloud are both virtualizable and nonvirtualizable live media production services such as ingest, vision mixing, multiviewers, and playout. The provisioning, construction, and control layers within the architecture discussed are all deployed within the private cloud and interact with the media production services in realtime. For resilience, services are deployed within two data centers called *Tech Hubs*, which are located approximately 400 km apart. Nothing within the architecture precludes systems and services from being deployed in more than two Tech Hubs. Great care has been taken to ensure that the systems within the Tech Hubs are loosely coupled, such that a system failure in one Tech Hub will not cause an outage in the other.

The next great opportunity is to take this to the public cloud where it can be coupled with a broader geographical reach and leverage the evolution of storage and compute. Logically, we can see the evolution and know it is where the industry needs to go. Transport formats that work well in the cloud today such as NDI and cloud digital interface (CDI) from Amazon use compression and bundled audio/video in a similar way to SMPTE 2022-22 formats are used. They are inherently simpler, take less bandwidth, and work more reliability in networks that are less defined or predictable. But they add latency, have limited audio tracks, and are less flexible. However, if the industry follows a familiar evolution we can hope for higher bandwidth, greater scale, and more flexibility in the future. Can we maintain the flexibility of ST 2110 in the cloud? Can we maintain the timing needs for live sports production in the cloud?

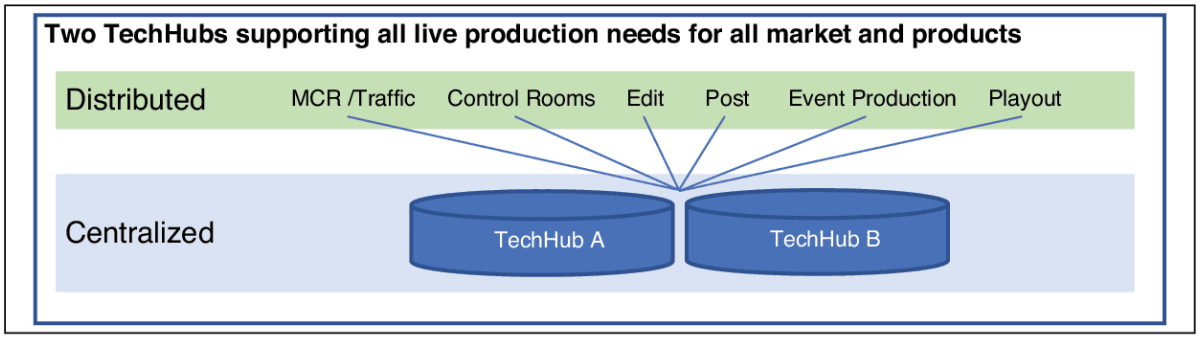

FIGURE 2. Distributed production with centralized storage.

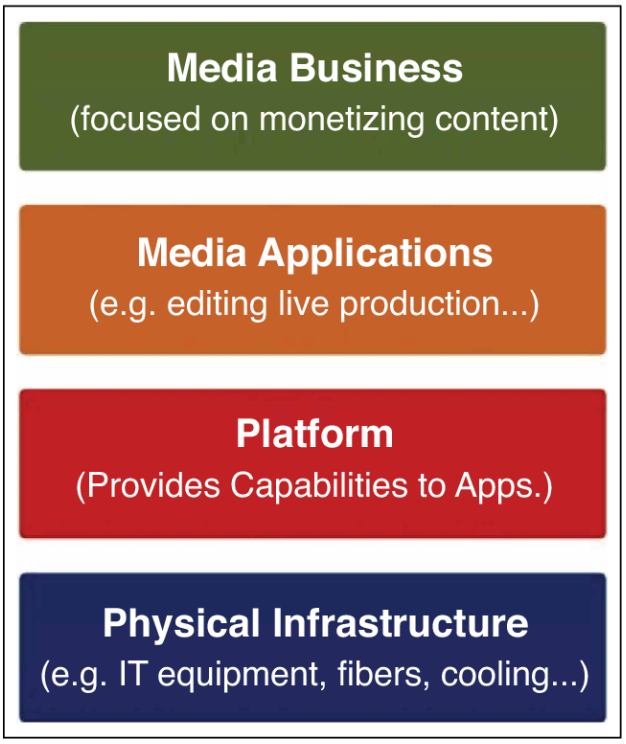

FIGURE 3. Generic media business architecture.

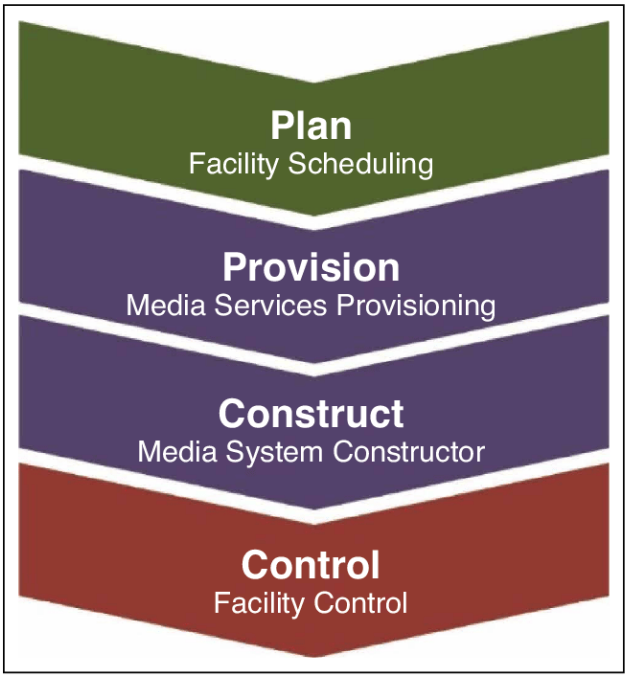

FIGURE 4. Dynamic architecture.

Conclusion and Recommendations

As we look to move our live production work to the cloud, hopefully, we can be more aggressive in addressing some of the questions outlined in this article. Focusing on transport is not enough. We need to couple this with recommended operating patterns and greater industry alignment. Moving forward, industry leverage IT tools and adapt as needed. It is easy to say that broadcasting or media is different or special, but that does not mean we need to develop our own tools for everything. We should set the bar for customization and unique system management tools very high and only use media-specific tools that are tied to creative aspects such as editing. The more we leverage the IT market, the more we can focus on true agility and microservice development. DevOps-based media applications will enable continued evolution and the cloud will provide the scale and flexibility we need to keep up with the changing consumer landscape.

About the Author

Gordon Castle, a vastly experienced technology professional, connects market dynamics to strategic plans and then drives execution that improves business performance and growth. This work has taken him to four countries and positioned him at the forefront of industry change and enabled him to succeed in roles ranging from start-ups to large-scale corporate strategy. Most recently, as senior vice president, technology and operations, Europe Middle East and Africa (EMEA) for Discovery, Paris, France and London, U.K., he was responsible for developing, defining, and implementing the long-term technology strategy of Eurosport, Discovery's leading sports brand. Before joining Discovery, he held senior positions with a telecommunications company, Ericsson, Stockholm, Sweden. As vice president, head of strategy development, media, Internet of Things (IoT), and applications, he developed strategic plans that enabled new growth areas. Then, as the vice president, the head of strategy development, portfolio, and capital and resource allocation, he oversaw the strategy and allocation of resources for the Ericsson portfolio. He has more than 30 years of experience working in the media industry, which has included roles such as director of entertainment, media, and communications with Pricewaterhouse, Coopers, Atlanta, GA; Senior Vice President (SVP) and the head of cable news network (CNN) Technology, Atlanta, responsible for the News Group's technology strategy, implementation, and training; as well as serving as an SVP—senior technology fellow with Turner Broadcasting, Atlanta. Throughout his career, he has helped his organizations embrace innovation by conceiving, planning, and executing large-scale initiatives that have changed organizations and the way people work. This vision contributed to his reputation as a leader in his field. An American by birth, he is based in London, U.K., where he lives with his wife and three daughters.