THE BROADCAST BRIDGE
Connecting IT to Broadcast

Orchestrating Resources For Large-Scale Events

A Themed Content Collection from The Broadcast Bridge

Themed Content Collection

Introduction

By Tony Orme. Editor at The Broadcast Bridge.

This collection of articles dives into the technology required to use IP and remote production to meet the potentially complex challenges of large-scale multi-site events.

SDI and AES have stood the test of time and have delivered reliability for the broadcast industry. However, viewer expectations are expanding, they want to watch what they want, where they want, and when they want. Static infrastructures such as SDI and AES do not lend themselves well to this, but the emerging IP technology has the potential to deliver on this, and more.

IP brings something quite new to broadcasting, something we've never seen before, and that is signal connectivity that is used in many other industries. Due to the massive data rates and low latencies needed by video, the type of transport stream (or link layer) we've used has always been limited to a synchronous serial system. SDI and AES are two examples of these. As other industries have ploughed more money into R&D than the broadcast industry could ever dream of, infrastructures such as IP, Ethernet, and fiber distribution have progressed exponentially.

The low latency and high data bandwidths now available to industries such as finance and medical have not only reached the requirements of broadcasters but in some cases even surpassed them. Compute resource and advanced storage methods provide more power now than ever. Whether this is available in on-prem datacenters or public cloud facilities, the compute capacity now open to broadcasters is greater than ever. The most significant change though has been in how we access it, that's through IP.

The real power of IP is that it is transport stream agnostic. With the appropriate interfaces IP packets flow seamlessly between many types of link layer. For example, it may start life being sent across a 10G Ethernet network, but then could easily be converted to HDLC at an ISP, and then converted to ADSL to be delivered to a viewer's home, and then WiFi for distribution within the home.

The combination of IP compliant link layer connectivity and massively impressive compute and storage resource with unprecedented low latency and high data throughput has delivered infrastructures that are delivering flexibility and scalability, in addition to the expected reliability. It's this combination that is allowing broadcasters to meet the viewers demands.

Before getting carried away with our new found utopian ideal, the type of compute, storage, and network resource needed to deliver video is orders of magnitude more expensive than a typical office-type IP infrastructure. To keep latencies low and data throughput high, the type of infrastructure needed is very high end. It's up there with nano-trading in the finance market and cutting-edge image processing seen in operating theaters. Consequently, it's not only expensive, but has a tendency to be complicated.

That said, with just a bit of knowledge an IP infrastructure can be made to transport video, audio, and metadata. But to take advantage of some of the more advanced features such as scaling and flexibility, much more knowledge must be acquired, especially as moving images have some interesting quirks.

It's also tempting to think that live video from a major sporting event can be easily transported across the internet to deliver broadcast quality video and audio to a centralized studio hub. Although this is possible, to improve reliability latency is increased to take into consideration dropped IP packets. And there is always a compromise between quality, reliability, and cost. Consequently, IT telco links are often employed to provide remote operations with the internet acting as a backup. Although this isn't as easy as it may first look, it does provide greater resilience, flexibility, and scalability than we are used to in the broadcast space.

One of the "free" hidden effects of IP networks is their implied full duplex operation. Being able to send and receive data to an outside broadcast opens the possibility of controlling remote equipment. This often manifests itself as remote camera control, sound console control and on-site monitoring. Traditionally, separate networks would have been needed if a broadcaster wanted to control remote equipment, but with IP, it often forms part of the network.

Monitoring IP networks is a new concept for broadcasters as we must establish reliable video and audio essence as well as a reliable link layer. SDI and AES are incredibly reliable most of the time, but the quality of the links for IP networks vary enormously, and this leads to the broadcaster needing to monitor the link layer.

Quality of Service (QoS) and Quality of Experience (QoE) work hand in hand to improve the viewer experience. The QoS may relate directly to the link layer, but it's only when it is combined with the QoE that the true measure of the viewers experience is defined. This is one of the areas where broadcast television is separated from the needs of traditional IT as some aspects of QoE may show the link to be high quality but could still have devastating effects on the video and audio. For example, latency may not be immediately obvious in a QoE measurement, but high latency can lead to a poor viewing experience.

IP networks have opened a whole new range of possibilities for broadcasters as the full extent of high compute power, scalable storage, and flexible networks is now available to distribute and process real-time broadcast video and audio. The next challenge we have is learning how to harness the potential of scalability to deliver programming for today's ever demanding viewers.

Contents

Planning Is Everything

By Paul Markham. The Broadcast Bridge.

An examination of how to plan & schedule resources to create resilient temporary multi-site broadcast production systems.

Major sporting events have long been an important feature of many broadcasters' schedules. Traditionally, outside broadcast vehicles would produce a single programme on-site at each venue, and this would subsequently be sent back to base via traditional telecoms or satellite infrastructure for broadcast distribution. Increasingly though, remote production is becoming more typical; the video and audio feeds from multiple sources on-site are all delivered to a production hub which can be hundreds or thousands of miles away. This approach reduces the headcount needed at each site, with all the associated travel and accommodation cost that implies.

Detail your vision

The diverse range of technical options open to us now means that we're no longer governed merely by the possible. It's essential to consider the available technologies not as an end in themselves, but as an array of tools to help you deliver a great viewer experience, and on the road to that, a great working experience for your teams. Costs will of course be a factor, but the process should begin with documenting a clear and detailed understanding of the production vision for each event and venue, who will need to be involved in delivering that to the defined standard, and what good and bad looks like from their perspective in terms of the tools they're given. You'll also need to consider "non-functional" requirements such as security and resilience. As the project proceeds you'll need to evolve the vision and make robust compromises and priority calls; the more understanding and engagement you have across the board, the better the outcomes you'll get. It can be difficult, but try to first define the vision in non-technical terms - it should be focused on people and business outcomes rather than on technology and solutions. The ***why*** rather than the how.

Can we use the cloud?

It might be tempting to think that, in 2022, we can do all of this with the cloud - the reality is that, although cloud has various roles it can play, we're not quite there in terms of facilitating remote live production for sporting events. The most critical issue is the consistency of latency between sources, which is currently challenging to guarantee in a cloud production environment. Audiences might tolerate various technical compromises on resolution and compression, but they're unlikely to tolerate seeing key moments, such as goals, multiple times in quick succession because of a variance in source latency. This is an area under widespread active development though, so it's a problem we can expect to be solved before too long.

Connectivity basics

Newer technologies are starting to shorten implementation timescales, but planning for a major event normally needs to begin at least a year in advance. In instances where many broadcasters from different markets will be seeking to cover the same event, a lot of the on-site infrastructure is typically managed by a host broadcaster. In other cases you'll need to be engaging directly with venues. Either way, early dialogue here is essential to understand what's possible, and what the costs are likely to be. Don't automatically assume you'll pass your main video and audio feeds back to your hub in the same way for every venue. Usually an IP stream will be appropriate, but you might need to consider a more traditional method like satellite, or there may be a role for internet or cloud.

Even if you choose native IP cameras you're probably going to need some encoding kit to compress the signals for distribution, and you may need to install an antenna for Precision Time Protocol (PTP). Then you'll likely need a rack in a co-location facility somewhere to link the connections from the venues to some trunk connectivity back to your hub. The host broadcaster or venue will have recommendations here; they may well have existing connectivity available. If they don't they'll be able to advise what local telcos have access to each site. Telco connectivity can take months, so early planning is critical. Also vital is doing the sums to understand exactly how much data you'll need to be concurrently sending. Give yourself some overhead here for flexibility - 20-30% if you can - as you really won't want to be hitting a ceiling when that crucial sporting moment happens.

Back at your hub you're going to need corresponding installations to receive your feeds, decode them as required, and present them to your production gallery. Whilst cloud and browser-based production software exists in 2022, it's not yet up to the standards you'll be working to for live coverage of a major event, so we're talking about a traditional suite, though it could be IP-based, using something like SMPTE 2110, and it could be control surfaces linked back to remote hardware. This kind of approach can facilitate flexibility and resilience, but it does add technical complexity. As ever, these decisions should be made in line with your defined vision and desired business outcomes, not for the sake of doing something appealing with technology.

Redundancy

For all but the most budget-conscious productions, one set of kit and one connectivity path won't be enough. You'll want to consider what happens in the event of failures. This probably means duplicating the connectivity all the way back to your hub via a second telco and second co-location site. Telcos should be able to advise on the actual routes of their fibres, which you need to be as diverse as possible. The most common cause of physical connectivity failure is construction work along the fiber route; the value of diverse provision is somewhat diminished if the two telcos have actually run their fibers along the same or parallel ducts, particularly in urban areas where construction activity is likely!

If you're sending the feeds around the world, public cloud providers and internet are worth investigating as cost-effective backup connectivity routes. Often some kind of pay-as-you-go option is available, this means there'll be some basic set-up costs, such as cross-connect with your co-location centres at each end, but the bulk of the transit cost you'll only have to pay if it's actually needed. As technology progresses in the coming years, it's reasonable to expect we'll see more feeds distributed in this kind of way.

More than just the action

Whether a host broadcaster is filming the action, or you're putting in the kit to do so yourself, you'll need to consider the other elements in your overall vision. Do you need to plan for a studio or a stand-up position at the venue? Will you have an on-site commentator? Will you be doing interviews with participants in the "mixed zone"? Are you expecting to provide off-site coverage at local landmarks or peripheral venues? How are you expecting to populate your on-screen graphics?

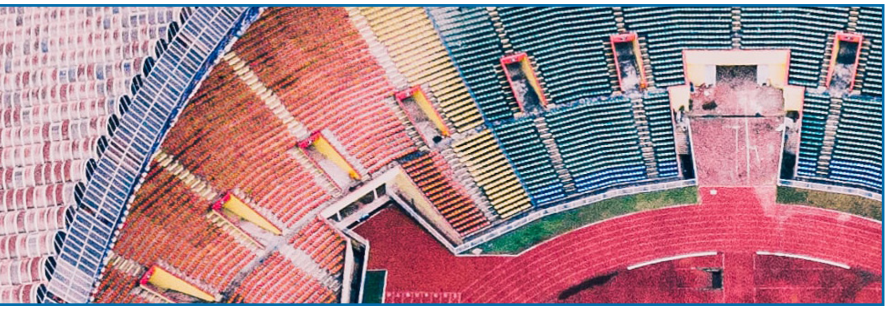

If you've defined these production aspirations in detail, you should be able to understand where latency and reliability may be less important and thus newer technologies such as 5G or cloud connectivity could be appropriate and cost effective. Backpack-type solutions are available for mobile reporting for example, and could connect over public 5G or 4G to your co-location centre or a public cloud deployment if needed, though of course it's important to remember that public cellular bandwidth is likely to be variable. By understanding these compromises in terms of business outcomes, you can make confident decisions about risks and ensure everyone knows what the limitations are.

Production logistics

If you have presenter talent at a venue you'll need logistics like confidence feeds, talkback, and presenter earpieces. And of course you'll need make-up, green room, welfare facilities, basic IT connectivity and technical support. Modern talkback and confidence feeds can easily be IP, and as most of your data is going to be travelling outbound from the venue to your hub, bandwidth in the other direction shouldn't be a problem. You'll need to consider what equipment you're going to need and how and where it'll be connected. Latency will likely be important for the audio communications but will be less important for video confidence feeds, so it's worth considering them independently. Again, your early engagement with the teams involved and understanding of their needs and priorities will ensure the best outcomes.

Network design and security

Modern networking equipment allows us to operate a single physical network but segment it into multiple virtual networks using Virtual Routing and Forwarding (VRF) at the IP layer. This provides the means to segregate the traffic from a security and quality-of-service (QoS) perspective. For our event we probably want to create at least three VRF segments (VRFs) - one for the main video feeds, another for control and logistics traffic including talkback and data for on-screen graphics, and a third for potentially "dirty" IT services such as email, telecoms and web browsing. This protects our video content from the risk of malware arriving via email, for example. We might also consider duplicating some of these so that we have software-layer network resilience to complement our physical network redundancy.

You should expect to perform some kind of security threat modelling or risk assessment against your overall architecture; for this, you'll need a detailed understanding of exactly what data is flowing where using what protocols, and how access is governed and secured.

Highlights and Archive

Alongside the live production, there's usually a requirement to store the material and produce highlights packages. This is where public or private cloud options can now come to the fore - no longer do we need expensive on-site edit suites - we can choose either fully-featured remote controlled desktop editing stations, or various browser-based Software-as-a-Service (SaaS) editing solutions.

There are some challenges though - public cloud is great for scalability, you can have capacity on tap, but you'll need to figure out how your video data is going to get there, how finished assets will be delivered to where you need them, what the data transit costs might be, and if additional data transfers might cause delays. You'll also need some cloud operations and security expertise to ensure everything works well and provides the best user experience for your teams. Private cloud deployments can be easier to get data in and out of, but you have to provision for your peak capacity requirement and do expensive upfront purchasing, installation and ongoing maintenance. Use your defined vision and business outcomes to decide what's best for your organisation, taking into account your existing infrastructure and expertise.

When it comes to archive storage of your material, public cloud is surely becoming the best option. Whichever provider you choose, long-term storage of data can be very cheap, flexible and site resilient. It's best to ensure that video is processed into a standard technical format for archive, considering frame rates, resolutions and audio track layouts as well as codecs and containers. Metadata should be stored alongside the video in flat files, JSON is currently the format of choice for this, and can allow for basic database-type operations and other integrations using tools within the cloud provider's platform.

Going live

Anything that can go wrong will go wrong, so an extensive testing phase is going to be required. Access to venues early enough to do this all in-situ is unlikely, so anticipate comprehensive configuration and testing of your venue kit at another location well in advance, potentially months ahead before you ship it from your base. At the other end, time will be short, so your on-site technicians will need detailed check-lists to ensure that the equipment has arrived undamaged and is installed correctly. If you can link these checklist items right back to your detailed vision and business outcomes, you'll know exactly what's impacted if anything does go wrong, and be able to make quick decisions about what to do or where to compromise.

Don't forget to schedule tests of your redundant paths and technical resilience as well - a useful approach can be to involve production and operations teams in "gaming" various resilience scenarios, so they know what to expect when the worst happens and can react accordingly.

Focus on Value

From the initial planning phase right through to decommissioning, put your audience and your people at the center of your decision making. As the project proceeds the teams will gain deeper technical understanding, which will necessarily inform an evolution of priorities. Equally external factors will spawn new requirements requiring robust assessment. Try to map every single piece of work back to a defined business outcome, and ensure that each outcome is properly understood by the team working to deliver it. The aim should be that teams are motivated by business objectives rather than technical aspirations or organizational tropes. This isn't as straightforward as it might sound - there'll be intricate dependency chains involved - but a clear understanding helps minimize the amount of wasted effort and delivers the most value.

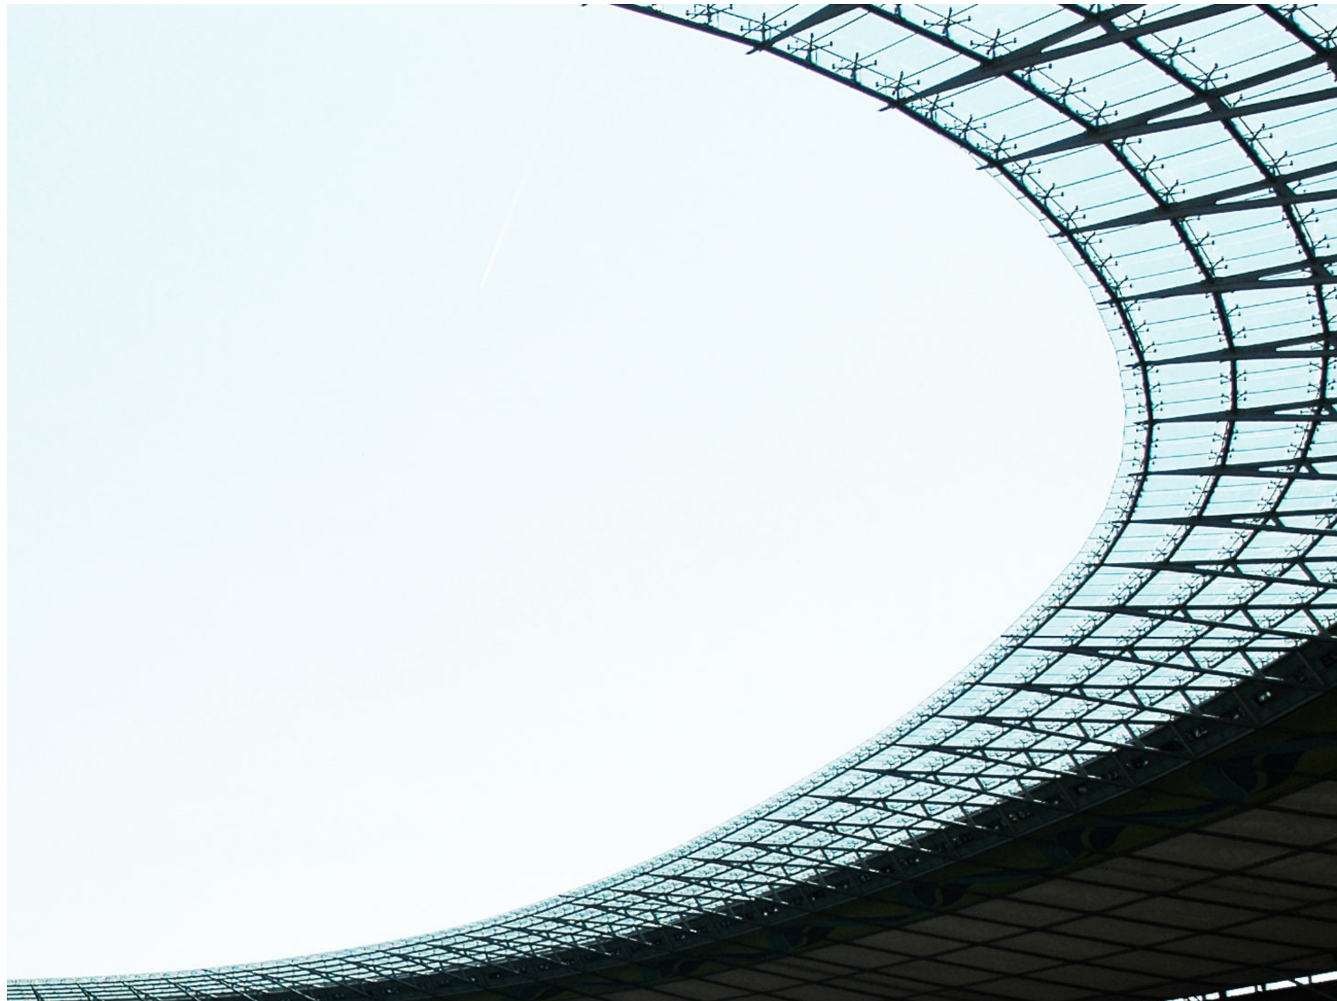
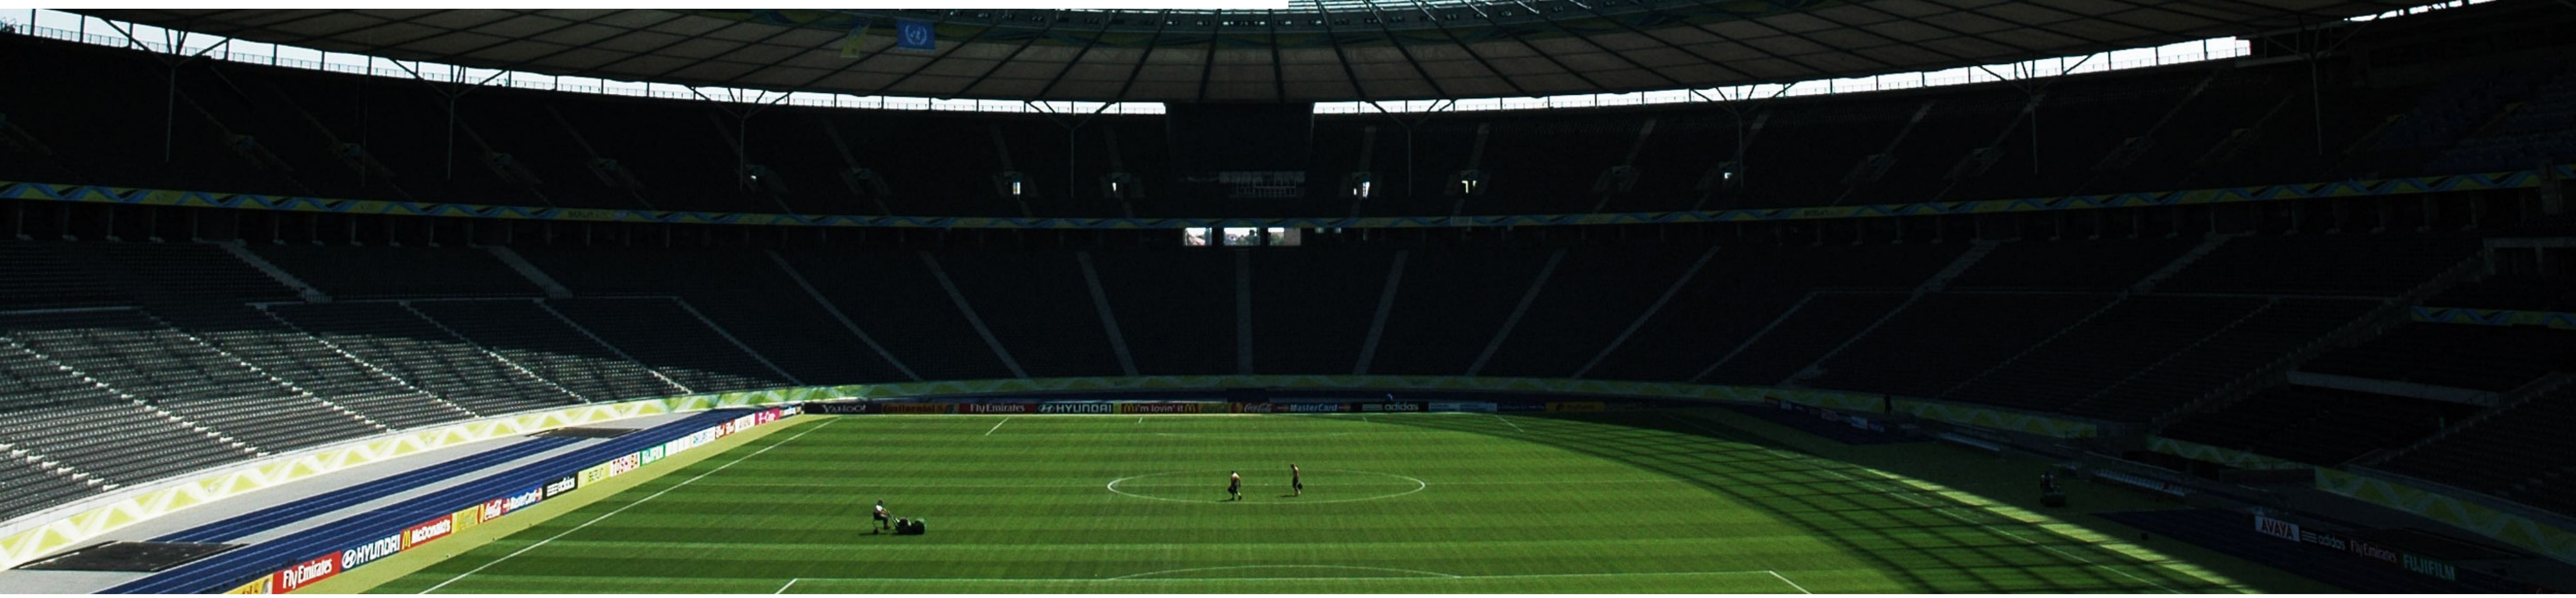

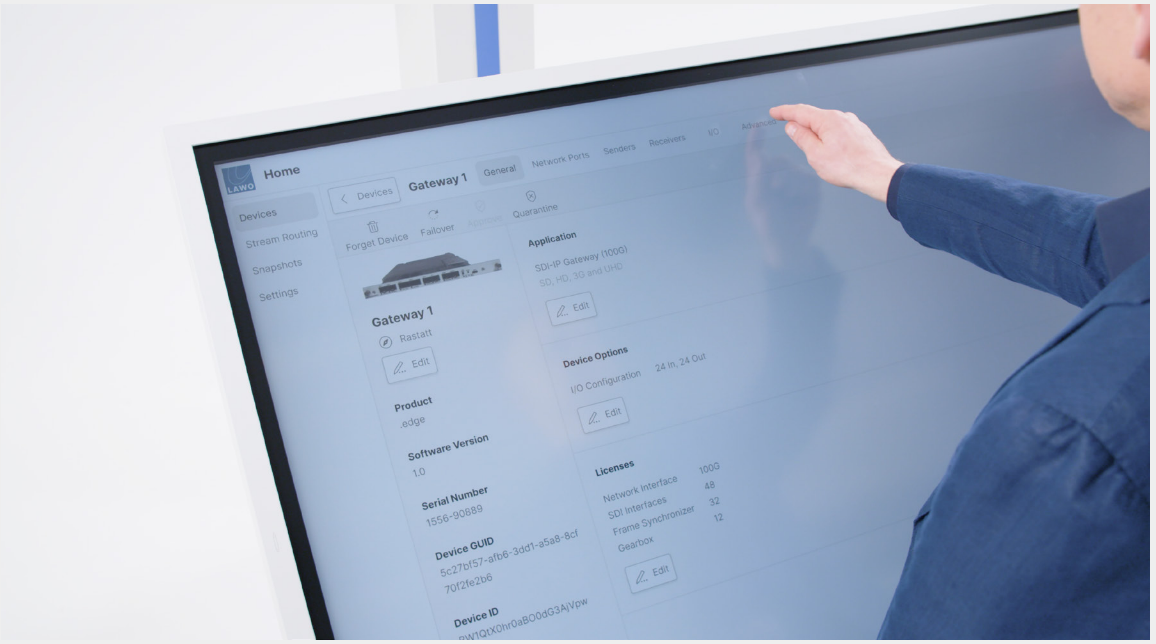

IP's Light-Bulb Moment

Axel Kern. Senior Director Cloud & Infrastructure Solutions at Lawo.

Simplifying the challenges of day-to-day operation with IP production systems using Lawo's HOME software environment.

In all fairness, open-standards-based broadcast IP is an excellent approach for turning clustered or distributed production sites into a unified network, and sharing compute and production resources scattered over locations that may be thousands of miles apart.

If you have already migrated to IP, you may, however, recall how complex the matter is behind the scenes. You need to be aware of IP addresses, subnets, and whatnots. And you depend on someone who is able to configure your infrastructure. The latter is often part of the package you purchase from the vendor or systems integrator. Still, the better part of the commissioned configuration only works reliably as long as no IP devices or processing cores are added or removed. Even simple things like disconnecting a device from one wall box and plugging it into another in a different room may require tweaks to the configuration that, more often than not, turns into yet another all-nighter.

Assuming that you have one or several IP-savvy engineers on your team, configuring a network for outside broadcast assignments still takes time. Time most of us no longer have, because of ever tighter production schedules.

Playing to The Galleries

'Old-school' configuration sessions of IP networks and devices involve assigning unique IP addresses to all devices as a first step, and then telling all other relevant devices where to look for the data they need to receive and/or process. The larger the infrastructure, the longer the preparation work—and the likelier you are to assign a given IP address twice, or to omit one figure.

In response to clamors from the field, several attempts have been made to simplify IP. NMOS IS-04 and IS-05 for easy device discovery and registration is one of them and is supported by a large number of vendors, mostly in the video field.

But didn't our customers actually ask for a system where "You connect everything, you push a button and it configures itself?" In an ideal world, this would include automated assignments of IP addresses with little or no human intervention.

One of the likeliest issues indeed used to be duplicate IP addresses. This is taken care of by HOME's built-in DHCP service, which automatically assigns IP addresses to devices and streams. On the whole, even a hyper-dense platform like .edge, can be up and running much quicker than ever before.

With HOME, this convenience has become a reality for video, audio and control devices. Not just Lawo products, by the way, because other manufacturers are encouraged to join Merging Technologies, DirectOut and others as HOME-native third-party devices. This status can be acquired by equipping or retrofitting devices with the HOME API. Doing so provides a number of advantages that go beyond discovery and registration. These are described below.

HOME covers all areas of the broadcast world: it speeds up initial infrastructure configuration, it makes remote productions easier to set up, and it saves operators of OB trucks massive amounts of time. This is because HOME has the ability to manage both local and remote resources, whether HOME-native or NMOS-compatible.

Although HOME can happily live in the cloud, most users will be relieved to learn that its containerized architecture works equally well on COTS servers in their on- or off-site data center. As a matter of fact, HOME is even available for live setups built around one mc² console and up to 20 stageboxes (from Lawo, Merging Technologies, DirectOut and NMOS IS-04/05 devices). Without requiring a dedicated license or COTS server, that is.

Why is that good news? Because connections are plug-and-play, without so much as a second spent on assigning IP addresses.

The recent addition of Global Snapshots, finally, allows operators to capture the current device and routing state of all audio and video devices in the system, and to recall those settings at a later stage. These Snapshots even include the routing states of NMOS-enabled gear.

What's in It for You?

We already alluded to the fact that HOME does not stop at automating device discovery, registration and IP address assignments. Another big plus it offers is a unified user experience based on the LUX (Lawo User Experience) design, which can be leveraged from the giant screens of an mc²-series console, or from any networked computer monitor.

It allows operators to change the settings on far-away edge devices, whether manufactured by Lawo or HOME-native third-party vendors. This makes it easy to tweak the input gain or phantom power status, for instance, without having to run

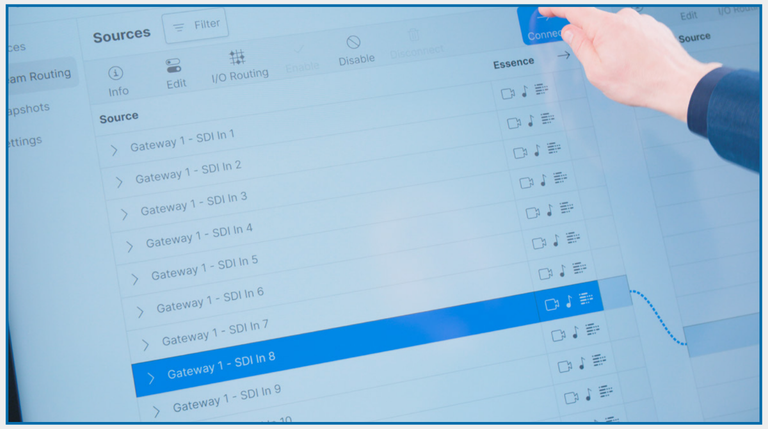

to the stagebox in question, or asking someone to change the settings for you.

In the case of .edge, all parameters are set within HOME, either as macro-like routing stream bundles (video, audio, metadata), or as individual essences for flexible audio shuffling and other deep-dive setting changes. In other words: the parameter complexity can be concealed and revealed as necessary. HOME acts as the official user interface for .edge, providing access to all orchestration aspects, parameter settings and configuration/system setup functions. Even software licenses for additional .edge features are managed within HOME. A similar integration is available for other (third-party) devices as well.

Thanks to its Signal Rights Management routine, HOME allows operators in a multi-console setup to “lock” the required I/Os of one or several edge devices to ensure that no-one else on the network can tamper with those settings.

The User List feature makes signal routing for mc² consoles in A__UHD Core systems more intuitive and flexible. Given the potentially high number of signals on networks dealing with large-scale productions operators can assign all required signals to logical groups that can be conveniently recalled at the press of a button.

Hardware monitoring is obviously included in the package: the System Health feature provides status information about a system’s components with centrally logged information for at-a-glance peace of mind during a production. Errors are categorized and displayed in real time according to their severity, making it easy for operators to decide on the urgency of taking corrective action. Optional DevOps tools like Prometheus and Grafana can be used for long-term logging and error analysis.

All in all, HOME lives up to the promise of its name : it is the go-to place for all things local or remote IP infrastructure, which can be conveniently managed and monitored in one place. Even NMOS IS-04- and IS-05-compatible devices can be routed, whilst parameter tweaking is not on the cards for them, because this is not included in the IS-04/05 specification. What is possible, however, is using HOME as a registry for both HOME-native and NMOS IS-04/05-compatible devices.

HOME can indeed retrieve information from NMOS devices through an IS-04 registry and supports connection management via IS-05 to provide at least some basic management functionality.

Is it Safe?

In an ideal world, no issues arise when you add more devices to a network. But we all know that real life works a little differently. Requiring operators to “admit” a device that is automatically discovered and registered for the first time is therefore a sign of good housekeeping: it gives you pause to double-check and correct settings that might conflict with other devices on the network and cause unforeseen collateral damage. This is exactly how HOME’s Quarantine process works. It may take a few seconds of your time, but that may very well save you hours of troubleshooting.

HOME Run

Designed to simplify managing IP infrastructures and save operators massive amounts of time, HOME is built on microservices and containerized in a Docker environment. This natively enables scalability detached from hardware constraints and is cloud-ready by design for future deployments.

Those who have worked with HOME, including broadcast service providers to the biggest planetary sporting events, agree that HOME makes their lives in an IP-based environment considerably easier. And that is still only the beginning: Lawo’s roadmap for HOME contains a few more milestones that will change your way of leveraging the power of IP-based infrastructures.

The content created by a production crew and transported over a network is your operation’s most valuable asset that deserves adequate protection, whatever its origin and route. While a robust security system needs to cover all aspects of media infrastructure and content creation, the key lies in its simplicity, its initial design and deployment as well as in its on-going maintenance and support.

HOME’s architecture is prepared for managing services such as transport layer security for user interfaces, control data and media essences, based on well-established IT security mechanisms, such as HTTPS, RADIUS, MACsec and IPsec.

Axel Kern. Senior Director Cloud & Infrastructure Solutions at Lawo.

Connecting Remote Locations

By Philip Hunter. The Broadcast Bridge.

Creating reliable, secure, high-bandwidth connectivity between multiple remote locations, your hub, and distributed production teams.

With a major international tournament where dozens of broadcasters around the world require access to content, the likely viable solution is distribution of shared feeds from a single set of cameras and audio sources to each broadcaster. A host broadcaster assumes local responsibility for camera provision and control, dictates the transmission format, and controls local telco provision. The broadcasters remote production hub may receive as many as 40 camera feeds etc from each stadium but probably won't have control over them.

When it comes to coverage of multiple venues for a national league on game day, the decisions revert to the broadcaster and their production providers, and that's what we discuss here.

Bandwidth On A Grand Scale

The amount of bandwidth required will dictate connectivity requirements, and that in turn depends on quality required and other components such as codecs. To transmit full uncompressed video requires a bit rate ranging from around 2 Gbps for 8-bit color depth with 4:2:2 chroma subsampling, up to as much as 8 Gbps for RGBA color at 16 bits per channel. 40 cameras for a high profile soccer game would therefore require between 80 Gbps and 320 Gbps of bandwidth. Our production hub systems might need to handle four stadia simultaneously. That sort of bandwidth is only available over direct fiber connections and compression will usually be required. This raises other issues, such as quality and also latency, since compression imposes a computational delay. When deployed in lossless mode to maintain quality, as may be required in remote production, the latency is greater still.

Intra codecs have evolved to reduce latency by compressing only within each frame to exploit spatial redundancy only, without taking advantage of temporal redundancy in the similarity between successive frames. This led to codecs such as JPEG 2000 and AVC-Intra, with recent focus on trimming latency to the bone for contribution of live events from the field, especially fast-moving sports. This led to development of the JPEG-XS codec, exploiting parallel processing to accelerate execution and with precise control over bit rate to enable remote use of available remote connectivity. JPEG-XS has enjoyed rapid adoption for remote contribution over the last few years as broadcasters have been attracted by the fact that it imposes very small delays in the order of a single frame. It offers compression ratios up to 10:1 for typical images, but can be higher in some cases, enabling use over professional video links in production scenarios that might have previously used uncompressed data.

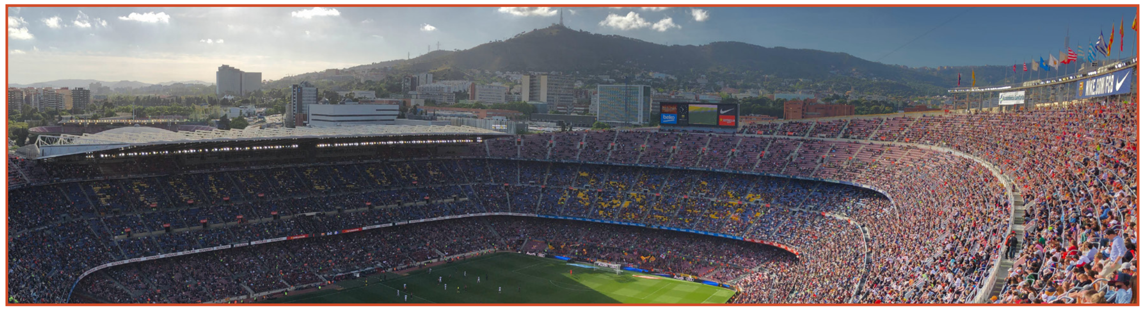

There are other considerations beyond bit rate and compression that may have a bearing on choice of connectivity method, primarily redundancy and error correction. Contribution flows compressed by codecs such as JPEG-XS can be carried inside MPEG Transport Streams, or increasingly SMPTE ST2110, which specifies carriage, synchronization, and description of elementary essence streams over IP for real-time production, as well as playout. A key benefit of SMPTE ST2110 is support for diverse paths to provide redundancy as insulation against packet loss and individual path failure within the IP fabric.

That is all very well, but such a fabric offering path redundancy may not be available in the field for remote contribution. In that case the only way of providing some protection at least against transient errors and varying transmission conditions, without imposing too much latency, is to employ FEC (Forward Error Correction). That inserts additional bits to provide a level of inbuilt redundancy within the data, such that IP packets can be recovered fully without loss in the event of some corruption, providing it is not too much.

All this has an impact on the decision over what form of connectivity to employ, if there is a choice. The use of FEC may require a slightly higher bit rate than the alternative of packet retransmission to achieve redundancy at the stream level, trading that for lower latency.

Planning, Control & Comms

Other considerations include choice of software tools for planning, configuring and managing the connectivity, beyond the codecs and handling of functions such as bonding. Various mixing and production functions are enabled by improved connectivity and lower latency, the most obvious being camera mixing. There is no standard approach to remote camera control in remote production; some prefer to keep control at the venue within an OB etc, others require control from the hub gallery, many are a hybrid. All approaches require comms between venue and gallery crew. In remote production connectivity then, it is increasingly hard to separate connectivity from the associated functions that depend upon it and in turn dictate the requirements.

Available Solutions

Broadcasters and content producers face an increasingly bewildering range of options for connecting their remote sites for distributed production, but the choices can be narrowed by considering the specific locations and requirements. It is increasingly a case of horses for courses with the sites that require connectivity ranging from modern well-equipped stadia to small remote settings beyond the reach of cellular connectivity, never mind fiber or other fixed line options.

Available options now extend way beyond the satellite connectivity that used to serve SNG vans which are now often regarded as the expensive nuclear option in the absence of any viable alternative, rather than the default.

For a long time, leased lines and managed services over the internet or even private networks have been available in some locations, but the big change quite recently has been the emergence of cellular communications in the 5G era as a serious option not just for uploading video at relatively low quality but for serious contribution from the field. This brings huge benefits in many scenarios, enabling high bit rate and low latency connectivity to be set up on demand at short notice in some remote areas where direct fiber is not available and where remote broadcasting is required only temporarily.

The optimum choice then depends on the facilities available, the quality required and of course budgetary factors. When broadcasting from a major modern stadium, connectivity will usually be provided on tap with direct fiber connections providing all the bandwidth needed. Smaller and older sports facilities might have leased lines available, which were once the mainstay of corporate communications but now a relatively low cost and low bandwidth option, although still benefiting from high levels of availability and guaranteed bit rate. Those properties are important for remote broadcasting where erratic connectivity cannot be tolerated, or at the very least requires a reliable backup.

Internet connectivity through broadband connections may well furnish that back up in many less well-endowed smaller remote facilities that may be quite new to remote broadcasting in the streaming era, apart perhaps from rare occasions. In sporting leagues that may happen when a small club is drawn against a big one in a knock-out competition after a good run.

Increasingly mobile networks are providing back up as well as primary remote connectivity at the same time. This can happen through use of networks from multiple service providers, in the hope that if one fails another will still be available. Then in normal times when all networks are running the available capacity may be aggregated through bonding or other means, so that in the event of failure the uplink bit rate might be reduced, but at least contribution can continue from the field.

Bonding was first deployed for remote contribution from the field as an alternative to satellite uplinks in the 4G era starting from around 2010. Then, as now, one caveat is that to bond a significant number of circuits requires taking out data plans with more than one mobile operator. However, there are providers of managed services that handle the bonding, and those relationships with cellular operators, on behalf of the end customer, acting as connectivity brokers effectively.

Future Possibilities

The role of cellular in remote broadcast connectivity is poised to increase further as 5G-NR (New Radio) rolls out, bringing higher capacities and bit rates (up to 10 Gbps Uplink), greater levels of availability measured in percentage uptime or sustained levels of performance, and lower latencies. NR refers to the completely revamped radio air interface developed for 5G bringing greater efficiency and high capacity over a given amount of RF spectrum. This ushered in 5G, but further advances are coming during the 5G era, notably Standalone (SA) operation where the core network behind the radio cells is in turn reshaped for higher bit rate and, importantly for some broadcast operations, lower latency. It is only when 5G SA is rolled out that the full promise of ultra-low latency operation on an end-to-end network basis, as opposed to just within the radio access layer, is delivered. This is relevant for some remote broadcasting functions of growing importance, such as switching between alternate cameras around a sports field.

Under 5G SA, another option is emerging, or will do so over the next few years, called network slicing. The idea is that the overall capacity of a 5G network is split up into multiple components, each of which can have different levels of service measured in average bit rate, consistency of this bit rate, and latency, which again can be within specified bounds. In this way a network can be shared more efficiently between multiple types of user varying in their need for QoS (Quality of Service), without over provisioning by giving too high a performance where it is not needed. This may avoid the need for bonding in some cases and allow broadcasters to specify slices flexibly, varying the precise parameters as requirements dictate. There may be times or situations where they need particularly low latencies. In one demonstration by Italian broadcaster RAI, a 5G slice was provisioned offering a guaranteed 60 Mbps uplink, leaving another 50 Mbps slice for lower priority services, such as casual internet browsing. There was also the option to bond these two slices together to yield 110 Mbps, but with less guarantee over the last 50 Mbps.

Contribution & Remote Control

By Dan Duffell. *The Broadcast Bridge.*

A discussion of camera sources, contribution network and remote control infrastructure required at the venue.

In many ways not much has changed with high-profile sports production; we have a collection of cameras inside a venue, connected to a gallery where cameras are remote controlled, sources switched, graphics added etc. With remote IP based production what has changed is that parts of that infrastructure are no longer parked outside the venue, they are hundreds or thousands of miles away. Fundamentally, most of the contribution challenges of remote production revolve around the latency introduced by the distances involved. Latency is the key issue which drives which bits of the infrastructure can sensibly be kept 'at home' in the hub and what needs to be on site.

By necessity therefore some of our infrastructure (beyond just cameras and microphones) and associated crew, needs to be on site. Vital systems for networking, compression, monitoring, remote control and comms all need to be somewhere near the action. How this is done varies; when the pandemic first accelerated adoption, much of the time a full OB vehicle would be sent to the venue with a subset of its systems actually used. This approach is still favored by some broadcasters and production companies as it offers an added layer of redundancy to the infrastructure. A few years on we have a range of approaches in play; Preparation of flypacks is becoming a very popular approach, because it gives an opportunity for individual rack and wider system configuration and testing back at base, and the flypacks are highly mobile making it practical to deploy them nationally or internationally. Some have built much smaller OB vehicles designed specifically for remote production. Some are installing permanent pitch side racks at larger venues where remote production is now a weekly occurrence on game day.

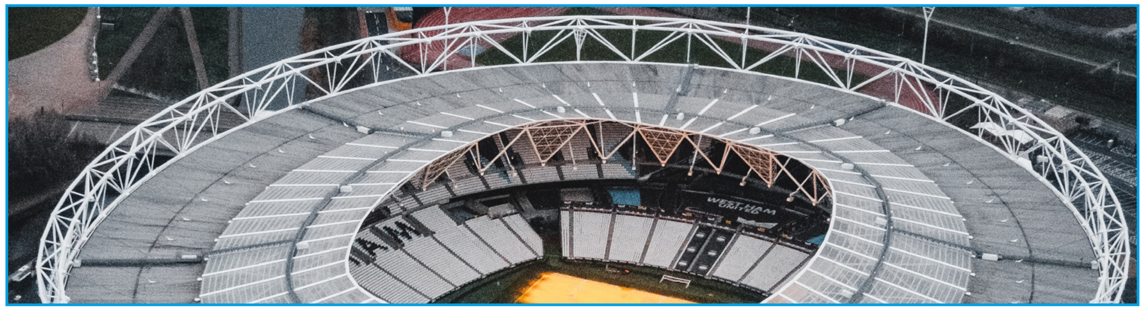

What are we connecting?

A typical high profile soccer game might have 40 cameras in and around the stadium. It's a combination of different camera types, each of which has its own requirements in terms of connectivity and control. Some are manned on site (system cameras, flycams, drones etc)… others are fixed position (stationary cameras, mini-cams etc), others are robotic (ptz, rail cams etc). Different cameras have varying degrees of remote-control capability. Most will have iris control for basic shading and many offer remote color control.

The rigging of cameras, audio etc is mostly pretty traditional and well understood. Connectivity to the contribution network is quite traditional too, either wireless or wired… serial or IP. Using IP broadcast cameras which are designed for this type of application will streamline connectivity and control. Serial broadcast cameras can be used with the aid of camera mounted converter boxes. Many specialty cameras will not be IP native and they too will need converters. It is possible to also combine contribution via cellular 4G or 5G, especially for roving cameras and positions away from the main venue.

Contribution Network

We are discussing remote IP based production so the backbone within the venue is an IP based contribution network. Essentially a high-capacity network switch that is taking input sources from all the cameras and feeding them to a transmission interface that connects to/from backhaul. Most prefer to use an uncompressed ST 2110 based contribution network, although some might opt for a compressed format, primarily to reduce cost by reducing required switch capacity.

Calculating required network bandwidth is a matter of mathematics: To transmit full uncompressed video requires a bit rate ranging from around 2 Gbps for 8-bit color depth with 4:2:2 chroma subsampling, up to as much as 8 Gbps for RGBA color at 16 bits per channel. 40 cameras for a high-profile soccer game might therefore require between 80 Gbps and 320 Gbps of bandwidth. This might work for an on-site contribution network, given sufficient switch capacity, but not for backhaul so some sort of compression before feeding into backhaul is almost certainly required. There are a number of codec options but JPEG XS is gaining in popularity because it delivers upwards of 10:1 compression without noticeable degradation in quality. Compression is typically achieved using hardware encoders located with your network switch(es) in your flypack, rack, OB etc. This compression inevitably adds some degree of computational latency.

Managing IP addresses is a perennial challenge in all large-scale IP systems. Every single device on the network needs to have a unique IP address. In addition to cameras, we must consider confidence monitoring, audio systems, comms etc. What may seem like a relatively manageable challenge within the confines of a single venue contribution network, becomes exponentially more complex within a multi-site IP network for a major tournament or national league coverage. It is a challenge that can be solved in a number of ways and one which is addressed by a number of commercial integrated platforms or IT systems. Regardless of the approach, a significant exercise in planning, device configuration and system testing is required ahead of the event.

Remote broadcast infrastructure requires a return path from the production gallery to the venue for confidence monitoring (video and audio for various members of the on-site team), GPIO, tally and remote control. Comms requires two-way communication. The venue contribution network needs to accommodate this in terms of design, configuration and capacity.

Remote Control

Most large-scale productions will be based on proprietary systems supplied by the major vendors. In this scenario, system cameras, CCU, routers and switchers are all engineered to work seamlessly together, and with IP native systems that means remote control is part of the infrastructure. With an IP native system camera, a single IP connection carries video, monitoring and remote-control data. Serial cameras are less straightforward; using a camera mounted converter box to bridge the gap between SDI and IP for video is simple enough. For remote control, Serial cameras present a latency related challenge; most will time out and revert to on board control (and settings) if there is too much latency between the camera and the CCU. Some converter box vendors solve this by running software within the converter box which sends data to the camera to fool the camera into seeing the converter box as the CCU.

Alongside the system cameras a significant array of specialty cameras is likely to be used and these are likely to require a third-party remote-control solution. Such solutions require a control data stream to be delivered to the camera via the network alongside video connectivity, and that may mean a dedicated IP connection.

There are broadly two different types of remote camera control to consider; shading/coloring, and motion control for lenses and/or robotics. What will be possible depends entirely upon the specific cameras and controllers used.

With shading and coloring, most cameras support iris control or if necessary, gain control. Many cameras support control of blacks and basic color. IP native system cameras may support more comprehensive remote color control. Many large-scale productions are utilizing automated color control systems which are part of the hub gallery infrastructure – significantly reducing the need for remote color control. With some remote-control systems, where cameras have an OSD (On Screen Display) it is possible to access this remotely, giving a fairly deep level of remote configuration capability.

Remote motion control presents a potential latency pain point. Once we are over the hurdle of delivering control data from the controller back in our hub, to a device in the venue, there is a simple operational challenge. How quickly does the operator need the device to react to commands to remain viable? The answer to that will be driven by production requirements and must be considered on a device-by-device basis. Control data requires significantly less bandwidth than video, so GPIO, tally and control data will potentially reach the device ahead of, for example a video or audio confidence monitoring stream, or indeed comms. During the planning phase, as the potential hub<>venue system latency becomes a known quantity, decisions can be made about where the operator for each type of device needs to be.

Audio

Most sports production relies on microphones in fixed positions. Using soccer as an example; there are typically 12 pitch side mic's. Stadium ambiance may be as many as 8 mics, with American football pitch side parabolic mics are used to capture the action for line outs etc and most matches will have wireless mics on the roving pitch side steadycams… but in all cases the connection of mics to stageboxes is traditional stuff. Any camera audio feeds are packaged with the video feed.

There is a requirement for an on-site audio engineer to rig, line check and do confidence monitoring so it is highly likely that a local touch-screen based console in a rack, with its own processor will be used. Such an approach would then typically provide capability for remote control over the entire on site mix engine (including mic' gain) from back at the hub using a console control surface. Having the audio processing engine on site reduces latency for monitoring.

With soccer, automated mix systems where video analysis is used to follow the ball around the pitch, opening/closing channels and setting levels for the nearest pitch side mic' is common. This would be operated back at the hub from the console control surface, sometimes using a second audio processing engine. Operators tend to be able to learn to live with managed latency.

Some Things Never Change

One thing has not changed; the success of any large-scale production depends on careful consideration, advanced planning, meticulous preparation and a lot of testing. What works for you will be determined by a careful examination of your own very specific requirements.

European Championships Produced With Flexible Comms Matrix Over Distributed IP Network

By Marc Schneider. *Executive Director Global Events, Riedel Communications.*

Riedel share their insider insight into how comms and network infrastructure was achieved for the massive 13 stadium 2022 production.

For production companies working on live sports production projects, the single most important technology that ensures a smooth outcome is crew communications. Without it, problems could occur and the director never knows about it until it's too late. There are always time-critical decisions to be executed and everyone involved needs to be in sync. Therefore, immediate and clear communication that addresses the right recipients, when they need it most, can be indispensable.

Getting it right requires an appropriately designed communication infrastructure that ties the various venues to each other by sending the right audio signals back though a main production center for processing and system monitoring.

For the 2022 European Championships this summer, all crew intercom communications were routed on a distributed network—with a MediorNet fiber-optic routing platform from Riedel at the center—that allowed virtually every crew member to communicate with any other person across multiple locations, both separately and simultaneously as a group. This IP infrastructure connected a total of thirteen venues in and around Germany: eight were connected over an existing dark fiber network within the city limits and five more at Olympiapark (site of the 1972 Summer Olympic Games), some 14 miles away.

The entire production was run remotely. At the 2022 European Championships it was operated from the IBC in Munich, but could also have been operated from Riedel's Remote Operations Center at its headquarters in Wuppertal, Germany. In fact, the entire installation was connected to the ROC for additional support and resources.

Distributed Fiber-Optic Network

More than 200 MediorNet frames were used, allowing the producers of the event—The Munich 2022 Local Organizing Committee—to send signals incoming from every venue to any output or even multiple outputs using a router control system. This gave the production team a great deal of flexibility in how and where comms signals were distributed and significantly reduced cabling and set-up time (requiring months of planning and about six weeks of installation).

"This fiber-based routing system eliminates the need for re-wiring if production configurations change at the last minute," said Felix Demandt, Senior Project Manager at Riedel. "In total, we laid over 1,600 km [about 1,000 miles] of fiber for this event."

The MediorNet installation was located inside the Technical Operation Center (TOC) in Munich. From there, the correct signals were supplied to the various international rights holders in the IBC as well as to the venues in the Olympiapark, and to other external venues.

"We received three 10 gigabit Internet in the Olympiapark and distributed it to the various areas such as Production, Media/Press, and Commentary," said Demandt. The IT backbone for communications alone consisted of over 600 switches. "In addition, we also provided WiFi with over 300 access points for production."

Software-Defined Orchestration

Events like the European Championships can benefit from SDN Orchestration to make the infrastructure more flexible and resilient, and easier to set up. To help the production go smoothly and to ensure that all signals (intercom, video, audio and data) are sent and received as desired), technicians can make use of real-time IP network orchestration and Software Defined Networks control systems like those developed by Belgium-based company SDNsquare, which Riedel acquired in September of this year.

This control layer can help teams customize infrastructures to accommodate various venues and seamlessly incorporate all of the video feeds, audio signals (including comms) and control data into a tailor-made and scalable networked production.

An IP network (based on MediorNet VirtU IP switches) was set up on site at the IBC and was used not only for access points at each venue but also helped integrate all of the comms activities as well. The entire comms infrastructure was based on the AES67 audio networking standard and was deployed using multicast routing, whereby signals were sent to a group of intercom users simultaneously.

At each venue the crew used a series of Riedel Artist intercom systems with Smartpanels as control interfaces and Riedel Bolero beltpacks, a wireless intercom capable of supporting up to 250 comms beltpacks and 100 antennas in a single deployment. An Artist-1024 talk back frame was also located at every venue, connected by AES67 to two frames in the IBC and an auxiliary frame in Wuppertal.

“We actually wanted to set up one intercom ring based on Riedel Artist in the Olympic Park, but we exceeded the dimensions of the system with 1024 subscribers, because we ended up with almost twice as many,” said Norbert Garske, Head of Broadcast Engineering Host/LOC “We therefore installed a broadcast matrix for the complete TV production and an event matrix, which Riedel managed themselves. The two rings were connected by trunk lines.”

[Of note: Up to 128 ports can be accommodated in a single frame and up to 128 frames can be connected to a redundant dual fiber-optic ring resulting in a maximum matrix size of 1,024 x 1,024 intercom ports.]

The Cost-Effective Freedom To Roam

Giving the crew the freedom to roam, the Bolero beltpacks support Bluetooth, which allowed the crew to use either a compatible headset or a smartphone. Crew members could receive calls on their phone and talk and listen via their headset. Users could also inject phone calls directly into the intercom channels.

Due to the massive size of the crew at all locations, the equipment complement also included the use of TETRA-compliant radios. Terrestrial Trunked Radio (TETRA) is a European standard for two-way transceiver specifications.

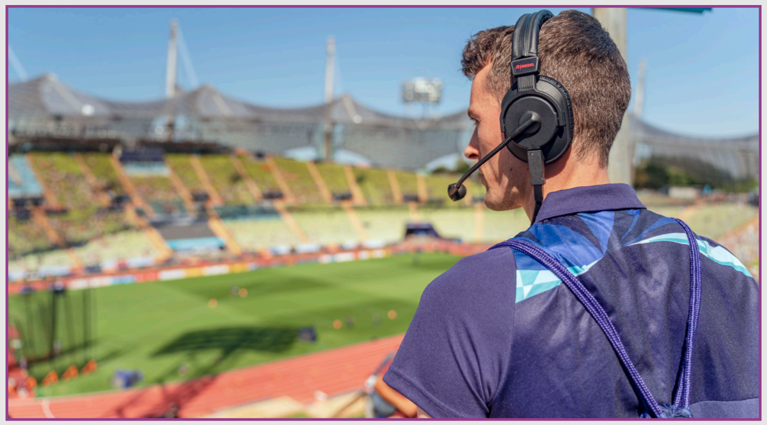

“With TETRA, you can generate very cost-effective coverage for personnel that do not need complex high-end solutions,” said Demandt, adding that Riedel provided 96 support employees. “Thousands of Boleros for people like site crew would mean unnecessarily high costs, so expanding the IP comms network with TETRA is a very efficient way of covering very large user groups.”

Ensuring QoS Comms

During ten days in August, the European Championships required 2,500 radios, 400 Bolero systems and 450 SmartPanels as a Web-connected user interface. Once a match was completed, the Local Organizing Committee then distributed it to broadcasters all over the world

Although the infrastructure was a bit complex, signal quality was foremost on the producers minds, as crew commands had to be heard correctly, with no static or dropped calls. To ensure QoS, the team deployed Quality of Service (QoS) DSCP Marking, which was used to determine which network traffic required higher bandwidth, has a higher priority, and was more likely to drop packets. They also made extensive use of Differentiated Services, a protocol for specifying and controlling network traffic by class so that certain types of traffic get precedence - for example, comms traffic, which requires a relatively uninterrupted flow of data, was given precedence over other kinds of traffic.

Comms Is Critical To (Sports Production) Success

Live event production is extremely challenging under the best of circumstances and you’ve only got one chance to get it right. In the world of live sports production, crew communications play a critical role across all production disciplines and are as important to the success of a major TV events like the European Championships as the performance of the football teams themselves.

Marc Schneider. Executive Director Global Events at Riedel Communications.

Leveraging Data For Monitoring, Efficiency & Forensics

By Ned Soseman. *The Broadcast Bridge.*

Large-scale remote production systems can be complex and challenging to monitor, but IP presents many opportunities to capture and make use of rich data streams.

The goal of every live TV sports broadcast is to provide the best possible content with the highest technical quality. How that goal is achieved usually depends on people, technology, and budgets. How do large-scale productions with strong technical talent monitor a production for compliance and efficiency to provide the highest QoS and QoE?

Throughout the history of live TV, technical crews have used waveform monitors, professional video monitors, and VU meters to visually confirm video and audio quality and metrics. Similar live signal monitoring continues to this day, but modern monitoring methods also include automatic logging, electronic alerting for technical errors and deviations, and capturing and logging everything for later evaluation and auditing.

Because the trend in major TV sports production is IP and REMI IP, digital performance metrics make it easy to capture, log and monitor data, to alert operators to exceptions, and to reliably reproduce exceptions for troubleshooting.

Many exception settings are programmable, subject to an acceptable minimum level of service. The level of acceptable exceptions often depends on the size of the audience and what technical issues viewers and sponsors will or won't tolerate.

Objective monitoring rather than subjective observation is infinitely better and more scientific than how it was done in the analog days.

The scale of production for a major international sporting event is epic compared to, for example, the scale of budgets and revenues for broadcasting a local fishing tournament. However, the technical needs and pitfalls at both ends of the scale are much the same. Large-scale productions have more moving parts, higher financial risks and are typically more sophisticated and complicated.

IP technology has changed the logistics and economics of TV production workflows and enabled distributed production, by allowing sharing and control between production technicians and equipment in multiple locations. When we have multiple stadiums spread across different cities, with multiple matches being played simultaneously, it is a massive undertaking for the production crews and infrastructure involved and necessitates an array of data collection points which all need to be devised, established, monitored and archived.

Integrated Platforms

Thankfully, this large-scale challenge has attracted considerable innovation in recent years and there is a growing portfolio of systems available to help.

There is more to using IP and the cloud in live sports productions than image processing, compression, and transport. Many major broadcast manufacturers have introduced integrated platforms under a variety of brand names and monikers that provide a wide range of IP infrastructure management solutions and controls specifically for broadcasters. Whilst the primary purpose of most of the tools in the orchestration technology stack is designing systems, and then managing how they are spun up and down, most of them also present opportunities for data gathering and monitoring.

What's In The Stack

Some integrated platforms qualify as software-defined networking (SDN) systems, where the management and control plane is separate from the data plane. The management and control plane typically allows a control screen customizable to prioritize and meet the needs in a particular facility. A SDN controller must support the local network architecture because of the wide variety of network switches and LAN topologies. One of the most common local networks in media facilities is monolithic, where all media edge devices are connected to the same switch. The problem with monolithic switches is that they aren't scalable. A scalable networking choice popular in media and data centers is spine and leaf.

The SDN controller establishes a connection between a source and a destination. It communicates with source devices, destination devices, the network, or in combination. A service orchestrator allows scheduling of resources and enables prediction and allocation of required future network capacity. It also adds in-depth monitoring and a time dimension to an SDN controller.

Many integrated platforms are customizable to match the unique needs of stations, production companies and OB scenarios. Many allow live monitoring and review of virtually all data points from each source and the inputs and outputs of all processing units, servers and connections in the workflow, including distribution.

Some integrated platforms are cloud-enabled to facilitate migration to the cloud, or partially cloud-based to take advantage of some cloud capabilities, or cloud-native applications conceived, designed and built to run entirely in a public cloud. Some are available as SaaS, others can be local software.

Not all integrated platforms are equally sophisticated.

Many platforms allow users to leverage existing SDI gear in a hybrid SDI/IP environment including UHD and HDR delivery and can also include centralized production processing, facility interconnect, and multi-resolution, multiviewer monitoring. Some systems also use AI and augmented reality (AR) technologies to enhance their power.

More versatile integrated platforms can replace individual spreadsheets such as crew and resource scheduling and management, asset management, automated media operations, and transmission management. Some platforms are equipment agnostic while some others are nearly all proprietary. Some integrated platforms require a great deal of bandwidth. In UHD systems, redundant 100 Gb data connections are not unusual.

Analytical Quality Control

TV QC includes media files before they air, and internal auditing of content and related data after it airs. Sometimes internal QC auditing is referred to as a forensic audit. Most TV stations capture the content they broadcast 24/7/365 to document that every frame and word of every spot aired, should a sponsor question it.

In terms of pre-air file analysis, analytical quality control (AQC) refers to processes and procedures designed to ensure that the results of file analysis are consistent, comparable, accurate and within specified limits of precision. While file-based QC has been in use for some years, live and post-air QC relies on logging and compliance data.

Analyzing network and system performance after an event can identify the root cause of what happened if something went wrong and be used to improve future efficiency. Some systems capture every frame of video, others capture one frame per second for visual verification. Some others capture every video transition. When combined with detailed data logging for deep diving into what did or didn't happen all three methods work well.

Stream Monitoring

In a fully managed system, all workflow signal metrics are monitored for exceptions and key performance indicators (KPIs). Some fully managed systems also provide security features such as dynamic port management, password checks, automatic software updates and AI anomaly and pattern detection to reveal unexpected signal flows, logins, and other unusual behaviors on the network.

Many monitoring systems rely on Web Real-Time Communication (WebRTC) to provide web browsers and mobile applications with real-time communication (RTC) via application programming interfaces (APIs). WebRTC allows audio and video communication to work inside web pages by enabling direct peer-to-peer communication, without plugins or downloaded native apps. WebRTC is supported by Apple, Google, Microsoft, Mozilla, and Opera.

Traditional stream monitoring begins with methods and toolsets to monitor content during the ingest/encoding/transcoding process of content prep. Incorrect levels during ingest will cause problems later. Stream monitoring usually also includes monitoring the core network during the delivery process, such as the origin server and CDNs, and by monitoring the player or app used to decode content in a typical end-user viewing scenario such as OTT or live ABR internet streaming.

IP Monitoring

Monitoring live IP transport streams typically includes detecting anomalies such as blockiness, frozen frames, black frames, and audio/video syntax errors. Loudness detection and correction are also monitored for compliance. Transport streams should also be monitored for QoS metrics. There are many solutions available for transport stream monitoring. Some solutions are designed to monitor ST2110 and ST2022-6 networks. Others are designed to monitor ABR MPEG-2 transport streams. Some can work with both.

In an adaptive streaming network, a single video program is transcoded into many bit rate variants the player can use to adjust for network performance. Bit rate "shifting" can sometimes occur while the content is being played. If the shifts are not smooth, playback anomalies can occur. All the bit rate versions must be properly aligned with each other for smooth bit rate shifts. Network conditions can change, which is what ABR is designed to compensate for.

Managing Complexity

There are a wide variety of choices for content and IP monitoring solutions. Facilities and production scenarios are unique. All solutions deserve due diligence to determine which is the best fit for your organization and situation.