



MXL IS THE FOUNDATION. HERE'S WHAT IT TAKES TO BUILD THE HOUSE.

A whitepaper by AWS partner TVU Networks

MXL is the Foundation. Here's What It Takes to Build the House.

A whitepaper by TVU Networks and Amazon Web Services

The Hard Part Isn't the Technology

Picture a college basketball game. Two teams, a few thousand fans in the arena, and a regional sports network that has promised its subscribers a live broadcast. The production team has done this before — they know the sport, they know the workflow, they know what good television looks like.

What they also know is that getting the show on air will take most of the day. Not because anything is broken. Not because the team isn't skilled. But because connecting a live production in a software-defined environment — routing camera feeds from the arena to cloud processing, configuring the application parameters for each software function, verifying that the encoder at the venue is talking correctly to the switcher in the cloud — is a process that requires a specialist engineer working methodically through a checklist of settings that would fill a spreadsheet. By the time the opening tip-off arrives, that engineer has spent six to eight hours on configuration work that has nothing to do with producing good television.

It's just not as simple as SDI.

This is the dominant reality of cloud live production today. And it is the problem the industry most urgently needs to solve.

This isn't one team's bad day. Primary research conducted across more than ten live production organizations confirms it is the norm. For a 90-minute production running on cloud infrastructure, the compute cost — the actual cost of the servers doing the work — runs to approximately \$15. The configuration and labor cost for setting up that same production runs to approximately \$890. The infrastructure is nearly free. The expertise required to assemble it is not.

That ratio — not the cost of compute, not the cost of software licenses, but the cost of configuration — is why cloud live production has been slower to take hold than the industry expected. The economics only work if you can reduce the setup burden dramatically. And that is exactly what a new generation of technologies, led by an open-source initiative called MXL, is beginning to make possible.

What MXL Is — and Why It Matters

For as long as broadcast television has existed, the industry has needed a common language for moving video between pieces of equipment. In the analog and early digital era, that language was SDI — a standard copper cable connection that any camera, router, or monitor could speak. SDI made the multi-vendor broadcast facility possible. You could buy a camera from one company, a router from another, and a monitor from a third, and they would all work together because they all spoke SDI.

When the industry moved to IP networks, a new common language was needed. SMPTE ST 2110 provided it — defining how broadcast-quality video travels over standard IP infrastructure, and enabling facilities to scale their signal routing far beyond what SDI's physical constraints allowed. ST 2110 was a genuine achievement, adopted by major broadcast networks and large production facilities worldwide.

But ST 2110 was designed for a world of dedicated hardware and managed networks. As the industry now shifts to software — where video doesn't travel over wires at all, but passes between applications running on the same server or cluster of servers — a new common language is needed again.

That language is MXL, the Media Exchange Layer.

The simplest way to understand what MXL does is through a comparison. The traditional way of passing video between software applications is like making a photocopier — every time one application needs to hand a video frame to another, it packages that frame, addresses it, sends it, and waits for the other side to receive and unpack it. Every step burns computing power. Every copy takes time. In a complex live production with dozens of software functions each processing hundreds of signal paths simultaneously, that overhead adds up fast.

MXL's approach is different. Instead of making copies, MXL gives every application a key to the room where the original video lives. A frame is written once into a shared area of the server's memory. Any application that needs it — a switcher, a graphics engine, a monitor — simply walks in and reads it directly. No copies. No packaging. No wasted computing power. The frame doesn't move; the applications come to it.

This matters for two reasons. First, it is dramatically more efficient — software-defined production workflows running on MXL require significantly less computing infrastructure than those relying on traditional data exchange methods, which has a direct impact on infrastructure cost. Second, and perhaps more importantly, MXL is an open-source project governed by the Linux Foundation, with contributors including the EBU, CBC/Radio-Canada, the BBC, Grass Valley, Riedel, Lawo, AWS, NVIDIA, and others. Any vendor can implement it at no cost. Any broadcaster can use it without locking themselves into a single vendor's ecosystem. In a market where vendor lock-in has been one of the biggest barriers to software adoption, that openness is not a minor detail — it is the point.

MXL is one layer within a broader industry vision called the EBU Dynamic Media Facility (DMF) — a reference architecture that defines a complete model for software-defined production. For more on the DMF, see tech.ebu.ch/dmf.

The Foundation Is Laid. What Still Needs to Be Built.

MXL is the right foundation. But returning to our production team at the basketball game — MXL alone does not get their show on air with less configuration work. It does not connect the encoder at the arena to the cloud. It does not help the engineer find and authenticate the camera feeds from the venue floor. It does not prove to the network's legal team that the footage is genuine and hasn't been manipulated.

Those problems require solutions that sit alongside and above MXL — solutions the industry is actively developing but that are not yet fully in place. Understanding what they are, and how they connect to MXL, is the difference between seeing MXL as a complete answer and seeing it as the beginning of one.

Getting signals in and out. MXL's shared memory model works within a single server or a local cluster of servers. It does not, by itself, extend to the arena across town or the remote venue across the country. Productions that span physical locations — which is most live sports — need a way to bring those remote signals into the shared memory environment efficiently and reliably, across the variable and unpredictable conditions of real-world wide-area networks. This is the connectivity layer that MXL depends on but does not provide.

It is worth pausing here to understand what happens when MXL meets the cloud — because the cloud changes the equation in ways that are not immediately obvious.

In a traditional on-premises environment, MXL's shared memory model operates within the physical boundaries of a server or a tightly coupled cluster. The network between applications is fast, predictable, and under the operator's control. In the cloud, those assumptions break down. Video flows may traverse multiple virtual network layers — from the hypervisor's virtual switch, through overlay networks that connect containers or virtual machines, across availability zones that are physically separate data centers, and potentially across regions separated by hundreds of miles. Each layer introduces latency, jitter, and the possibility of packet loss — none of which are tolerable in a live production workflow where a single dropped frame is visible to millions of viewers.

What makes cloud infrastructure compelling despite these challenges is that the raw compute is extraordinarily affordable and elastic. Unlike dedicated hardware, cloud capacity can be provisioned in minutes and released the moment the production wraps. There is no capital expenditure, no truck roll, no hardware sitting idle between events.

The challenge, then, is not whether the cloud is powerful enough or affordable enough for live production — it clearly is. The challenge is bridging the gap between MXL's local shared

memory model and the distributed, multi-hop reality of cloud networking, without sacrificing the simplicity and efficiency that make MXL valuable in the first place.

Making sources easy to find and connect. Even with MXL in place, connecting to a remote camera or an encoder at a venue requires knowing where it is, how to reach it securely, and how to configure the connection correctly. Today that process is manual — and it is where most of that \$890 in configuration labor goes. What the industry needs is a mechanism as simple as pairing a Bluetooth device: the encoder at the arena registers itself, the production team finds it in a list, and the connection is made. Emerging specifications like VSF TR-12 are working toward exactly that, and leading vendors are building practical implementations now.

Proving that what you're seeing is real. As AI-generated video becomes indistinguishable from genuine footage, the ability to cryptographically prove the authenticity of a live signal is becoming essential — particularly for news and sports organizations whose credibility depends on it. The C2PA initiative is developing standards for content provenance, currently focused on distribution formats. Extending that to live contribution and production — to MXL flows and other live video formats — is an important unsolved problem that the industry needs to address.

These are not distant challenges. They are the specific gaps between MXL as a foundation and software-defined live production as a practical, economically viable reality.

What a Complete Solution Looks Like

Connect the pieces, and the picture comes into focus. A complete solution for software-defined live production has four components working together:

A **common media exchange layer** — provided by MXL — that lets software components from any vendor share video efficiently and without lock-in.

Global connectivity that extends the shared memory model beyond a single facility or cloud environment, bringing remote venues and distributed production into the same efficient workflow.

Plug-and-play device discovery that makes connecting a production as simple as finding a source in a list and clicking connect — not manually configuring IP addresses and port mappings in a spreadsheet.

Content authenticity built into the contribution and production chain, so that organizations can prove the provenance of their live footage from the moment it is captured.

When these four layers work together, the engineer at the basketball game doesn't spend six hours on setup. They spend thirty minutes. The economics of cloud live production change fundamentally — not because compute got cheaper, but because the configuration burden that dominated the cost equation has been addressed. And the production team that stayed on SDI

because software was too hard now has a path forward that doesn't require a specialist with deep knowledge of every vendor's implementation.

That is the goal. It is within reach. And the industry is moving toward it faster than many realize.

TVU Networks and the Road Ahead

TVU Networks has spent more than two decades building live video infrastructure for conditions where failure is not an option — breaking news situations, remote venues, locations where network conditions are unpredictable and the signal still has to arrive. That heritage shapes how TVU has approached the gaps described above.

TVU MediaMesh® — TVU's unified platform for live media connectivity, processing, and delivery — extends the same transmission technology that TVU built its reputation on across the entire live media chain, from field contribution through production to distribution, in a single orchestrated platform.

On the connectivity challenge, MediaMesh addresses the LAN boundary directly through what TVU calls global shared memory. To extend the analogy used earlier: MXL gives every application a key to the room where the original media lives. TVU MediaMesh makes that room as big as the world. A camera at an arena, a graphics system in the cloud, and a production switcher at a remote facility can all participate in the same shared memory workflow — with the remote source appearing to the software process as simply as opening a shared file, regardless of where in the world it originates. MediaMesh is built to work with MXL, not to replace it — extending MXL's model to the distributed environments that MXL alone cannot reach.

On device discovery and connection toil, TVU's global object model — a core component of MediaMesh global shared memory — assigns a unique identifier to every frame of every video source in the TVU ecosystem, making any source addressable from anywhere without manual configuration. TVU is also an active participant in the VSF TR-12 Client Device Discovery effort, an emerging open specification that brings plug-and-play device connectivity to production equipment at the edge of the workflow.

On content authenticity, TVU has joined the C2PA initiative to contribute toward extending cryptographic content provenance to live video formats, including MXL flows.

MXL-based integrations with partner vendors — including partnerships being announced around the time of this paper's publication — reflect TVU's belief that the complete solution described above will be built by an ecosystem, not a single company.

Conclusion

The basketball game gets produced. The engineer goes home at a reasonable hour. The regional sports network delivers its broadcast without needing a specialist engineer on standby for every event. That outcome — which is entirely achievable with the technologies now being developed — is what the industry is working toward.

What makes this moment different from previous waves of broadcast technology transition is the convergence of two powerful cost advantages that have never existed simultaneously before.

The first is the affordability of cloud infrastructure. At approximately \$15 of compute for a 90-minute production, the cost of processing power is no longer a barrier — it is a rounding error. Cloud infrastructure is elastic, globally available, and requires no capital expenditure. The hardware problem, for all practical purposes, has been solved.

The second is the simplification of human labor. When MXL, plug-and-play discovery, and global connectivity work together, the configuration burden that currently dominates production costs — the \$890 that dwarfs the \$15 of compute — collapses. Setup that once took 6–8 hours is reduced to under 30 minutes, cutting costs from \$890 to roughly \$56.

Together, these two advantages do not just make existing productions cheaper. They make entirely new productions possible. The Tier 3 and Tier 4 sporting events that were never economically viable to produce live — the college conference games, the regional tournaments, the niche sports with passionate but smaller audiences — suddenly have a path to live broadcast. The economics that once reserved live production for marquee events now extend to any event where there is an audience willing to watch.

And the implications extend beyond sports. The same cost structure that unlocks lower-tier live sports also opens the door to live production in professional AV, corporate events, houses of worship, education, and use cases the industry can barely imagine today. Anywhere there is a need to produce, switch, and distribute live video — and the budget has historically made it impractical — this combination of affordable cloud compute and simplified production workflows changes the calculus.

MXL is the foundation. It is the right foundation, built the right way, with the right coalition behind it. What gets built on top of it — the connectivity, the discovery, the authenticity — will determine how quickly software-defined live production moves from promise to reality.

But there is a larger possibility here. MXL has the potential to become what SDI was for a previous generation: the universal standard that every piece of production equipment speaks, that every vendor implements, and that every production team can rely on. Not just for broadcast — but for every segment of the live media industry. The house is being built. The foundation is sound. And the neighborhood is bigger than anyone expected.

For a deeper technical exploration of the complete solution architecture described in this paper, see the companion deep-dive whitepaper: [URL placeholder]

For more information about TVU Networks and TVU MediaMesh, visit tvunetworks.com

About the Authors

Mike Cronk is Vice President of Strategy for TVU Networks, where he leads the company's strategic marketing initiatives to deepen TVU's customer-focused innovation and industry collaboration. Mike brings a wealth of experience from his leadership roles, including founding Chairman of the Board at the Alliance for IP Media Solutions (AIMS), Head of Product for Live Media Services at AWS, and numerous leadership roles at Grass Valley including Vice President of Core Technologies, Senior Vice President of Strategic Marketing, and General Manager of the Server & News Production Business Unit.

Cronk holds an MSEE and BSEE from MIT, is the author of multiple patents, and holds a Technical Team Emmy award for his contributions to NBC's broadcast of the 1996 Olympic Games.

Simone D'Antone is a Global Strategy Leader for Broadcast at Amazon Web Services (AWS), where he focuses on media and entertainment workloads with a specialization in live production and contribution. Simone works with broadcasters, sports organizations, and technology partners to design cloud-based architectures for real-time media workflows. He is an active contributor to industry initiatives including the EBU Dynamic Media Facility (DMF) and has been instrumental in bridging the gap between traditional broadcast engineering and cloud-native production.