

The IP Convergence Framework for Distribution

How broadcast and streaming came together to address the C-band transition

Innovation rarely happens in a vacuum. More often, it happens at the point of collision: when worlds that have run on separate tracks begin to converge.

That is the story behind this framework. It began as a project for a broadcaster with a global distribution footprint of over 1,000 receive sites, whose ambition was moving away from a fixed satellite architecture toward a flexible, IP-based distribution network. To get there, the project brought together vendors from different corners of the industry that had never collaborated this closely in a modern, end-to-end workflow.

As the pieces came together, what this broadcaster and its vendors had built started to look less like a one-off solution and more like a framework: a repeatable model, adaptable across many organizations and contexts, even beyond broadcast.

And sometimes, timing lines up. The 2027 C-band spectrum auction is a hard deadline. Every broadcaster still relying on C-band satellite is evaluating alternatives before 2027 budgets lock. This framework happens to provide a working answer to this problem.

What sets this framework apart from other C-band approaches is that it moves master control to the cloud as part of a complete workflow, from ingest through distribution. This enables centralized operations, remote visibility, faster deployment, geographic flexibility, resiliency and disaster recovery, reduced dependency on physical facilities, and the ability to manage both streaming and traditional distribution from a common operational environment. The customer is unnamed, but the architecture is real and in production.

In the next sections, we will break down every element of the IP Convergence Framework for Distribution: where it came from, the layers that make it up, how they work together, and what benefits they deliver. We will also make the case for why this model is a working answer to the C-band transition US broadcasters are now facing.

The origins of the framework

Convergence between worlds

For most of broadcast history, “broadcast” and “streaming” have been treated as two different sectors, run by different teams, with different mindsets.

Broadcast was built around precision and visibility. Broadcast engineers typically think in frame-accurate timing and redundant signal paths. And, more importantly, they demand operational visibility: the ability to watch every piece of content moving and to intervene when something goes wrong. It is a discipline built the way Swiss watches are built, particularly in live broadcast, where every moment matters.

Streaming was built around reach and convenience. Streaming platforms were designed to serve media content to many different devices, locations, and bandwidth conditions, all at the touch of a button. If broadcast is like a Swiss watch, streaming resembles a modern logistics network: optimized for delivery at scale.

For a long time, these two disciplines stayed separate because the underlying technology was different. Broadcast ran on dedicated signal infrastructure, while streaming ran over the public internet. In recent years, IP has been closing the gap, enabling the best of both worlds to come together. That convergence is what gives this framework its name.



Figure 1: IP is the foundation bringing together the worlds of broadcast and streaming.

Convergence, in action

This convergence between streaming and broadcast, enabled by IP, was the ambition at the heart of the project that gave birth to this framework.

The broadcaster behind the project relied on a fixed satellite distribution architecture that was, by design, expensive and inflexible. This broadcaster has a global footprint of over 1,000 receive sites, including remote and hard-to-reach areas requiring reliable delivery. Reaching every region meant building and maintaining dedicated infrastructure locally, market by market. While this made sense a few years ago, it scaled poorly

against how content is consumed today, with requirements for more formats and devices moving faster than physical infrastructure can keep up with.

The broadcaster wanted to transition to an IP-based, software-defined architecture that combined the reach and convenience of streaming with the precision and visibility of broadcast.

The relocation of the broadcaster's facilities forced the decision. Rather than rebuilding its satellite plant at the new site, the broadcaster used the move to build its distribution on IP from the ground up, cutting costs in the process.

The new system was built around the following key design principles:

- **Software-defined:** Infrastructure abstracted from physical hardware and moved to the cloud, including master control, replacing fixed local deployments with elastic, centrally managed resources.
- **Flexible:** A cost model built around optimized variable spend rather than fixed provisioned capacity, cutting costs compared to the previous setup.
- **Content-agnostic:** Non-live content delivered through a streaming-first workflow, with live content receiving the control required by broadcast workflows.
- **Resilient:** Redundancy built by design, deployed across two cloud regions, protecting against signal loss.
- **Transparent:** Full visibility into every piece of content moving through the system, with the ability to intervene for operational staff (if needed).




Figure 2: The transition from fixed satellite to modern IP architecture.

“There’s no better example of embracing the idea that broadcast and streaming are two sides of the same coin – different faces, but ultimately part of the same thing”

Michael Atkin, President & CEO, BroadView Software

No single vendor could deliver all these dimensions at once. Most are strong in one or two. Broadcasters increasingly avoid closed platforms that force these tradeoffs and choose interoperable components from several vendors that work as one system.

The architecture of the framework

The layers

Four vendors came together to build the core foundation of this architecture:

- **Uplynk – the streaming and integration layer:** The layer handling streaming end to end, including encoding, editing and VOD publication. As the solution integrator, it also connects the other layers into one working system.
- **Zixi – the transport and reliability layer:** The layer moving streams into and out of the cloud, managing global resilience while also providing real-time network and stream health end to end.
- **BroadView + Nexus – the management and orchestration layer:** BroadView is the system of record for what is supposed to air, when and where. Nexus, its orchestrator, turns that schedule into instructions for the other layers, with real-time visibility and control across broadcast and streaming paths.
- **Grass Valley – the broadcast layer:** The layer handling broadcast-specific workflows such as distributing live sports content.




Figure 3: The four core layers of the architecture.

“It can give broadcasters a single solution that isn’t complex. An end-to-end solution built on best-of-breed technologies”

Steve Hassan, VP of Payout, Grass Valley

AWS is the backbone underneath all four: providing reliable global cloud resources on which to host the solution.

This framework is modular. Because the layers are interoperable, each piece can be swapped or extended to achieve something slightly different or adapt to specific requirements. It applies to any organization that needs to move live or non-live content reliably to many locations at once, not just broadcasters, but increasingly enterprises too. In fact, it has already been adapted to a second deployment.

In the project that originated this framework, some point solutions were layered on top of the core vendors mentioned above, including:

- **TAG:** For continuous multi-viewer monitoring and analytics across signals.
- **Videon:** For backup ingest encoding (satellite to cloud redundancy).
- **Unity Intercom:** For communication between control rooms.

Together, these layers covered every requirement the broadcaster had at the start. The next section looks at the workflow they enabled.

The workflow

The layers above describe who does what. The workflow describes how content moves through them, end to end, to deliver on all the key principles the project set out at the start: software-defined, flexible, content-agnostic, resilient and transparent.

The figure below illustrates the workflow graphically, in a simplified fashion.




Figure 4: The end-to-end workflow diagram.

These are the key steps of the workflow, in more detail:

- **Content ingest:** Feeds reach the broadcaster via satellite at an existing reception facility, or, increasingly, via SRT direct to the cloud, with Videon providing a backup ingest path when a primary feed is unavailable. Whatever the source, every feed is received in the Zixi platform in the cloud. Satellite signals are carried into AWS over the Zixi protocol from the reception facility, joining the SRT feeds, so everything is managed and monitored in one place.
- **Content lands in a virtual router:** Inside AWS, Zixi Broadcaster acts as a virtual router, doing the job a physical video router does in a physical facility, routing every incoming signal to one or several destinations.
- **Live content routes to cloud control rooms:** For anything airing live, the router feeds one of several cloud-based master control rooms (MCRs), built on Grass Valley's AMPP platform and designed to feel familiar to operators used to a physical control room. TAG's multi-viewer tiles run continuously across these rooms, giving operators a live view of the signal and quality analytics for every feed passing through.

- **Everything else routes to Uplynk for recording:** The same router simultaneously feeds Uplynk's cloud transcoders, which record incoming content in HLS at a fixed 10 Mbps, regardless of the bitrate it arrived at, to control storage costs without compromising the quality required at delivery. This is one of the areas where convergence becomes most practical. HLS is a streaming protocol, built to serve consumer devices, but here its top layer is encoded at the broadcaster's chosen distribution quality. The result is a single recording that meets the needs of both consumer and broadcast endpoints, with no separate broadcast-grade copy to store, track or manage. It is a streaming format doing a broadcast job.
- **BroadView decides what gets recorded and when:** As the system of record, BroadView knows the schedule. Where record lists once had to be manually keyed into the previous, on-premises playout system, BroadView's Nexus orchestrator now generates them automatically, starts and stops each recording, and leverages the Zixi platform's stream switching capabilities to route inputs to the correct encoder or control room. Nexus also provides a broadcast-style routing panel so that operators have the same familiar control capabilities they would expect in a physical plant. Automation is the default rather than the only mode: Nexus breaks each task into the steps that execute it, shows operators those steps as they run, and keeps a record of what ran and what failed.
- **Editors prepare content before it airs:** Recorded content becomes available in Uplynk's Clipping Studio, a non-linear editor built for HLS content, where operators trim, mark ad breaks and confirm content quality.
- **The path splits by content type:** Streaming-first (e.g., VOD, FAST) content plays out directly from Uplynk via its own channel scheduler. Content destined for live playout is instead exported as an MP4, deposited in an AWS S3 storage bucket, and picked up by Grass Valley for its playlist.
- **Nexus optimizes costs:** Across both paths, Nexus tracks each asset's lifecycle, and archives content once it is no longer needed, so storage does not quietly accumulate the way it would in a purely manual workflow.
- **Content reaches 1,000+ endpoints (2,000 planned):** Zixi multiplexes each of the broadcaster's outbound channels (16) and distributes them, leveraging the Zixi protocol's enhanced recovery mechanisms and network optimizations where possible with SRT available to meet existing hardware requirements. Delivery is not tied to a medium either: satellite, LEO, cellular or others, depending on what each endpoint has available.
- **The entire workflow runs redundantly:** Two AWS Virtual Private Clouds (VPCs), deployed in separate regions, operate as primary and backup with ZEN Master handling automatic failover between them, so a regional AWS issue does not interrupt delivery. Unity Intercom keeps control rooms in communication

regardless of where staff is physically located. Zixi's ZEN Master gives operators end-to-end telemetry across every hop, including packet loss and packet recovery on each path, so a fault can be isolated to a specific network segment or vendor without guesswork.

“We needed to make it operator friendly. We wanted the operators in the new facility to not have to conceive a whole new idea of how things work”

Ray Festa, Director of Solutions Engineering, Uplynk

The impact of the framework

Benefits of this approach

The approach advocated by this framework offers several benefits to broadcast and media organizations, including:

- **Cost:** Savings are realized throughout the workflow. Recording at consistent 10 Mbps instead of preserving each feed at its original, often much higher, bitrate cuts local storage costs significantly, while the shift to cloud replaces fixed spend with elastic resource provisioning. Moreover, automated lifecycle management in Nexus enables storage and AWS resource optimization to further reduce costs, and switching endpoints to native Zixi protocol reception can cut cloud egress spend by roughly 20–30% due to the protocol's compute optimizations in tandem with lossless null packet suppression at the network layer.
- **Resilience:** The workflow is resilient at every layer and step. Two AWS regions with automatic failover orchestrated via the Zixi platform, multiple ingest paths per feed (including the Videon backup), with available hitless failover mechanisms designed for 99.999% availability (Zixi), transport that adapts to endpoint requirements (e.g., satellite, LEO, cellular, etc.) and advanced monitoring capabilities provided throughout the workflow (by TAG, Zixi and Nexus). All at sub-second latency.
- **Location independence:** Because MCRs are virtual, operations can run from the main facility, from a home environment or from a mobile device, with no loss of function during disruptive events that would strand a traditional facility. This is an added layer of resilience on top of those mentioned earlier.
- **Operational trust:** The workflow was designed to mirror how a physical facility works, so that the operators using it are not learning a new mental model on top of a new architecture. On top of that, operators are given the same visibility and

control they would expect from a traditional facility: Nexus is where that control is exercised, while multiple monitoring options, from TAG's multi-viewers to Zixi ZEN Master's telemetry, provide visibility into stream health, packet loss and recovery across the workflow, so faults are pinpointed rather than searched for.

- **Flexible operating model:** Broadcast-specific production control (Grass Valley) and fast, streaming-first workflows (Uplynk) run on the same underlying infrastructure, coordinated by a single schedule of record in BroadView and Nexus, so broadcasters are not forced into a single, inflexible operating model. The same flexibility extends to scale: new endpoints, channels, or FAST channels can be added with a single click and delivered over whatever the destination requires, from any industry-standard transport protocol to alternative paths like Ku-band or cellular, all without redesigning the architecture.



Figure 5: The main benefits of the IP Convergence Framework for Distribution.

“That’s the beauty of this deployment: we don’t need to roll it all out day one. We scale as they scale”

Phil Abbott, Principal Solutions Architect, Zixi

Why this is a working answer to C-band retirement

In July 2026, the FCC voted to auction 160 MHz of upper C-band spectrum, with the auction required to close by July 2027. As Hive Group’s previous research on this topic noted, this is a structural disruption to the satellite-based distribution infrastructure that has underpinned contribution and delivery across the United States for decades.

Broadcasters are preparing to shift their satellite feeds to alternative distribution methods, including Ku-band, managed IP and hybrid IP, among others, ahead of the 2027 deadline. Chosen as a direct replacement for C-band, each comes with tradeoffs. Ku-band, for example, replicates the C-band operating model but is vulnerable to rain fade. This architecture removes the need to pick one. Each endpoint can be served over the path it requires from the same architecture, with terrestrial recovery available to protect Ku-band paths against rain fade.

The parallels to the broadcaster behind this framework are hard to miss. In fact, it faced similar challenges: a distribution model that had become too expensive and inflexible, a hard deadline (the facility relocation) forcing the decision, a need to preserve broadcast-specific workflows, and a need for reliable delivery to hard-to-reach areas around the world.

While this solution was built to solve a specific problem, it turned out to be a working answer to a broader one, for the following key reasons:

- **Flexibility by design:** Broadcasters can request only the feeds they need, migrate upstream providers one at a time rather than all at once, access control rooms from any location, and scale endpoints without redesigning the architecture.
- **Broadcast-specific workflows and performance survive the move to IP:** Latency stays at the broadcast level, live content gets the full production control, and operators keep the visibility and intervention points they are used to.
- **Built-in resilience protects broadcasters from disruptions:** Redundancy at every step of the workflow, including backup ingest paths, dual-region failover, and continuous monitoring.
- **The cost reduction is already realized:** This architecture has driven significant cost savings compared to traditional satellite distribution methods, including a substantial drop in local storage in the project behind this framework.

Conclusion

Innovation rarely happens in a vacuum. And sometimes, it happens serendipitously.

The partners behind this framework set out to solve one broadcaster's migration problems and ended up with a model that any media company at the convergence between broadcast and streaming can rely on to operate at that intersection, rather than choosing one world over the other.

MCR in the cloud is one of the clearest examples of these two worlds coming together. The virtual router with a broadcast-style panel is another. They marry the precision of broadcast with the flexibility of IP technology. And IP is ultimately what has enabled

these companies to make this framework function flexibly, reliably and efficiently. Flexibility from swappable layers and scalability. Reliability from redundancy and continuous monitoring. Efficiency from storage and transport cost optimization.

This framework has laid out its origins, architecture, and, most importantly, its impact. And that is worth highlighting further. This framework describes something already live in production for a real broadcaster, now scaling its distribution toward full deployment. It is not just theoretical architecture. And the benefits this approach has already achieved also happen to be a working answer to the challenges posed by the upcoming C-band retirement.

Sometimes you solve tomorrow's problems by working on today's.

To see the framework in detail, join the partners for a live webinar on Thursday, November 12 at 11 a.m. ET, where they'll discuss how this framework came together.

[Register here.](#)
