



The **Hybrid** **Playbook** for Broadcast

How to Optimize On-prem and Cloud
for More Agile and Efficient Operations

Version 1.0

Date of Publication

January 2026

Developed by

Imagine Communications

www.imaginecommunications.com

The Hybrid Playbook for Broadcast

As media organizations face increasing pressure to operate with greater agility, efficiency, and resilience, many are evaluating the role of cloud in their broadcast and production workflows.



This playbook explores

- where cloud delivers the greatest value;
- how hybrid on-prem/cloud models enable the balancing of cost, control, and performance;
- how software-defined workflows bring added efficiency to hybrid operations;
- and which technical, operational, and business factors broadcasters must consider when building and optimizing their own hybrid operations.

CONTENTS

- 03 Why Hybrid Matters
- 04 Cloud Economics Explained
- 06 Where On-Prem Excels
- 07 Operational Factors Beyond Cost
- 08 Utilization and Cost Efficiency
- 09 Designing a Balanced Hybrid Model
- 11 Cloud-Optimized Workloads
- 12 On-Prem-Optimized Workloads
- 13 Orchestration Across Environments
- 14 Hybrid Workflows in Practice
- 15 Cloud-Based Disaster Recovery
- 16 Scaling for Live Events
- 17 Efficiency and Environmental Impact
- 18 The Future of Hybrid Operations
- 19 Summary



On-Prem + Cloud: Greater Agility, Efficiency, and Resilience

The conversation about broadcast infrastructure has moved beyond a simple “on-prem or cloud” debate, and for good reason. Cloud and on-premises are not competing ideologies. Both can be successful. The real advantage comes from understanding when each model makes sense for your business and using them deliberately, not dogmatically.

By implementing a hybrid approach that leverages the distinct and complementary advantages of on-prem and cloud, broadcasters can better balance cost, control, and performance, achieving cost optimization across the entire content lifecycle. When these environments work together, broadcasters gain the resilience and adaptability to keep pace with shifting audiences, technologies, and markets. They can turn infrastructure from a fixed cost into a source of efficiency and strategic advantage, ensuring resources are right-sized, workflows are optimized, and operations can evolve without disruption.

If it’s an economic fit for a broadcaster’s costs and OpEx model, compute can be situated in the cloud. Or, for organizations that prefer to spend capital, it can be physically located on-prem, possibly in a private cloud.

Regardless of where the infrastructure is deployed, the future lies in software-defined payout and live production, unified through orchestration. Orchestration enables flexible deployment — **on-prem, in the cloud, and in hybrid models incorporating both.**

The greatest benefit of hybrid operations comes with identifying the use cases that leverage the cloud most efficiently.

Optimizing efficiency requires balancing cost, latency, and performance.

Broadcasters can’t optimize for all three at once, but they can choose the right two for each use case.

On-Prem

Hybrid

Cloud

- Control
- Low Latency
- Predictability
- Always-on Workloads



Orchestrated Operations
Workloads Move Dynamically
Cost / Performance Balancing



- Containerized
- Loosely Coupled
- Dynamically Scalable
- Automated
- Resilient

Workloads shift based on utilization, cost, and operational needs

When Flexibility Meets the Bottom Line

Seeing the cost of lift-and-shift implementations, many broadcasters still view the cloud as being expensive. But well-architected cloud-native platforms often can support more cost-efficient operations than on-prem infrastructure. The economic advantage of the cloud depends on how much and how often broadcasters use those resources. At lower utilization, cloud wins through efficiency and agility.

Instead of purchasing and maintaining physical hardware, broadcasters can buy cloud capacity by the hour, gigabyte, or stream. This consumption-based model turns CapEx into OpEx, offering immediate scalability and reduced upfront cost. It yields elastic capacity that enables broadcasters to scale on demand and launch services faster.


More flexible than traditional on-prem models, the cloud opens up new options for remote production. Because cloud generated feeds can be distributed globally broadcasters can provide remote access to operation teams anywhere.

In addition to increasing reach with real-time access from multiple locations, this model brings added security and the agility to experiment with intent for faster time to market. When experimenting or innovating, broadcasters can fail fast and move on.

That said, with added flexibility comes added complexity. Each service, be it compute, storage, networking, or data egress, has its own rate card.

Add orchestration, monitoring, and managed databases, and the total cost can increase quickly. At the same time, increased complexity demands upskilling of resources. While cloud brings many engineering benefits, it also requires intentional investment in human resources.

When broadcasters understand these costs and benefits, and where on-prem and cloud reach cost parity, they can move forward in optimizing overall resource allocation.



The Cloud, Defined

The cloud is compute, storage, and networking delivered as elastic, metered services, offering capacity you can spin up when needed and shut down when idle. Public cloud delivers this capacity as a rented platform — effectively, someone else's computer.

Private cloud delivers this capacity on infrastructure you own but operate with the same software-defined approach used in the public cloud.

Both models enable agility, shared resources, and rapid scaling, making them ideal for low-utilization workloads, disaster recovery, and temporary capacity.

Software-Defined Workflows

Software-defined refers to broadcast functions that run as software on general-purpose compute instead of being tied to fixed, purpose-built hardware. In a software-defined environment, capabilities such as playout, encoding, switching, or processing are virtualized, portable, and orchestrated independently of the physical machines underneath them.

This model separates function from hardware, allowing workloads to move between different environments — on-prem servers, private cloud clusters, or public cloud instances — without redesigning the workflow. It also enables dynamic scaling, automated deployment, centralized management, and more efficient use of shared compute resources.

In practical terms, software-defined systems make the broadcast chain more agile and resilient: channels can be launched or updated quickly, workloads can shift to wherever capacity is available, and operations can evolve without forklift changes to infrastructure.

The Value (and Limits) of Owned Infrastructure

Broadcasters have perfected the art of building reliable, high-performance on-premises systems, which remain essential for always-on services, regulatory compliance, and low-latency operations. Owning the hardware means complete control, predictable latency, and tight integration with facility workflows.

When utilization is high and requirements are constant, on-prem infrastructure often delivers the best long-term economics. (Notable early adopters of cloud technology and workflows have recognized this dynamic and brought select services back to their on-prem facilities.) Yet, the full cost of “owning your own” goes beyond the server invoice.

Every rack carries overhead: power, cooling, real estate, maintenance, and staffing. Capital investments must be amortized over years, and equipment risks becoming a stranded asset when business priorities shift. There are also intangible costs: time-consuming design stages, slower commissioning and deployment, lack of agility, and limited scalability when demand spikes.

Knowing where on-prem and cloud reach cost parity, broadcasters can make decisions on where resources should live — financially, operationally, and architecturally — and where they bring the best value for the business.

The Full Cost of On-Prem



Visible & Fixed Costs

Server hardware

Storage arrays

Networking

Maintenance contracts

Software licenses

Support services

Hidden & Variable Costs

Staff time and expertise

Power & HVAC overhead

Facility space & real estate

Risk of stranded assets

Lost agility / slower deployment

Lifecycle replacement costs



Beyond the Spreadsheet: Managing the Intangibles

Some factors that influence technical and operational decisions simply can't be captured in a cost model. In practice, the true value of on-prem or cloud often depends on how much risk is tolerable—and not just the "go wrong" outage risk. Broadcasters also face risk due to changing geopolitical pressures, regulatory changes, and supply chain uncertainties.

On-prem deployment of infrastructure gives broadcasters control over their data and content, eliminating geopolitical risk but remaining susceptible to local environmental hazards and to supply chain failures. Cloud platforms remove some of that risk but introduce dependency on third-party infrastructure, as well as potential vulnerability to sovereign cloud requirements, public broadcaster

regulations, or local regulatory bodies that may restrict or prohibit public cloud usage.

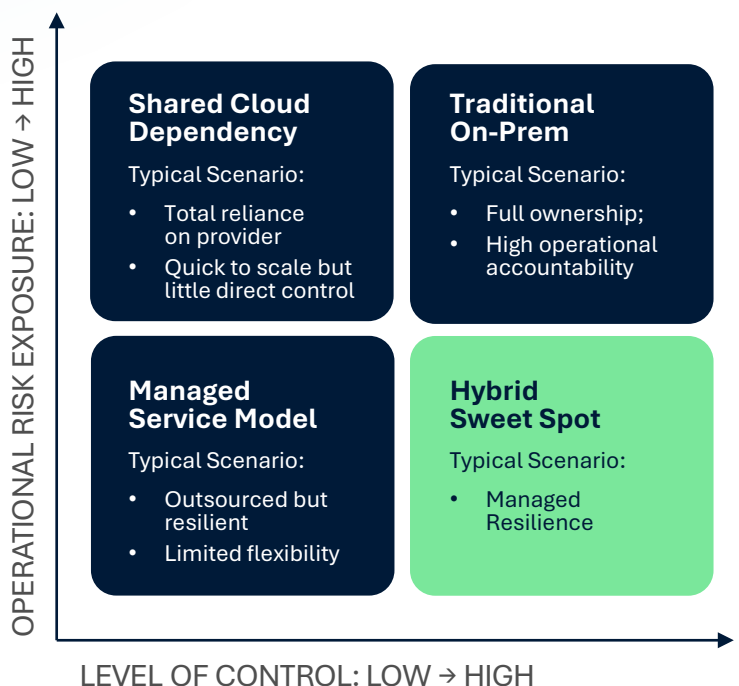
Whether deployed on-prem or in the cloud, broadcast systems are always at risk. The key to managing that risk lies in solid architecture and design, irrespective of deployment model.

Well-designed hybrid architectures enable broadcasters to address ever-present risks, along with today's intangibles. In maintaining hands-on confidence where it matters most, and by strategically offloading routine, redundant, or less-critical services to the cloud, broadcasters can minimize downtime, optimize recovery, and build resilience into the core of their operations.

The Software-Defined Skills Gap

Broadcast engineering teams have become expert at sailing on the winds of change — from analog to digital, SD to HD, SDI to ST 2110/IP, and, of course, the long road from tape-based to file-based workflows.

The latest shift in direction — from box-level products to software-defined workflows — presents another area in which engineering teams will need to build skills and develop working practices to apply new technologies to the task at hand: making and delivering great television.



Determining the Cloud-Ground Balance

When it comes to cost, utilization is the metric that matters most. The economics of cloud versus on-prem depend less on technology and more on how intensively each workload runs. Idle hardware wastes capital; over-spun cloud resources drain budgets just as quickly.

Imagine Communications' analysis of real-world deployments shows that cost parity occurs around 30% utilization. Below that threshold, cloud's pay-as-you-go flexibility wins the economics debate. Above it, on-prem infrastructure typically delivers a lower total cost of ownership and greater predictability.

Real-world examples make the pattern clear. Disaster-recovery systems may run only 6% of the time. Regional-split channels average 12%. Remote-edit suites reach 20%. Each favors the cloud. By contrast, 24/7 playout operates at 100% utilization, firmly an on-prem workload when a broadcaster is also managing failover, redundancy, and power costs.

Knowing where each workflow falls on the utilization spectrum is the foundation of any intelligent hybrid strategy.

This understanding empowers broadcasters to architect cloud infrastructure so that they use only the necessary resources, only for the necessary amount of time.

Failing to do so can be costly.

24/7 Playout

Resource-intensive
99.999% reliability target
24x7 operation

100% Utilization

On-Prem

Remote Editing

Cloud connected
99.95% reliability target
6 hours per day, 6 days a week

21.5% Utilization

Cloud Candidate

Regional Split Playout

Resource-intensive
99.999% reliability target
3 hours per day operation

12.5% Utilization

Cloud Candidate

DR backup for Playout

Resource-intensive
99.9% reliability target
10 hours per month operation

1.5% Utilization

Cloud

Hybrid Optimization: The Smart Middle Ground

True efficiency lies in balance.

Neither cloud nor on-prem is universally better; each excels under different conditions. The most successful broadcasters design hybrid systems that combine on-prem reliability with cloud elasticity, dynamically shifting their workloads to optimize cost, performance, and operational resilience. Using orchestration to establish a consistent operational control plane across on-prem and cloud solutions, broadcasters further reduce operational cost while realizing workflow savings.

Traditional broadcast redundancy, such as mirrored A/B systems, delivers high availability but doubles hardware cost and idle resources. Hybrid architectures replace that fixed duplication with intelligent, task-specific resiliency, such as using cloud only during a planned maintenance window, or standing up additional regional feeds for part of the day. Orchestration also enables asymmetrical N+M protection, where a shared pool of backup instances can support multiple primaries based on business value and fault tolerance.

The result is resilient operations without permanent overprovisioning, lowering cost while improving flexibility and responsiveness.

Hybrid also enables a smarter form of control.

Orchestration and automation across cloud and on-prem environments provide the foundational framework for consistent operations. This enables a unified operator experience, allowing operations to scale up for live events and scale down during periods of lower demand hours, and maintain continuity even when one environment fails. Decisions are guided not just by technical requirements, but by cost, utilization, and business value.

The outcome: lower risk, higher efficiency, improved sustainability, and an infrastructure that remains adaptable as business needs evolve.



Four questions to ask in considering hybrid for your business

1 ✓

Which workloads would benefit from shifting between on-prem and cloud?

Not every workflow needs elasticity. Identify the ones that do.

2 ✓

When and how often do your capacity needs change — by hour, by region, or by event?

Implement a hybrid model so you can scale only when and where demand exists.

3 ✓

What level of operational control do you need, and where can automation take over?

Streamline your operations without adding complexity for operators.

4 ✓

Are your architecture decisions guided by utilization and business value?

Capex requires long-term consistency of purpose, while elastic scaling allows risk to be managed over time.

Flexibility, Agility, and Scale on Demand

Cloud technology shines when workloads are unpredictable or intermittent. In these scenarios, pay-as-you-go pricing eliminates stranded capacity and reduces upfront investment.

Standing up new playout channels, training environments, and testing configurations — or handling short-term spikes, such as sports events, breaking news, or seasonal campaigns — can require just hours, not weeks.

Elastic compute resources, combined with orchestration tools, let teams scale operations instantly across regions or time zones. Environments can exist temporarily and then disappear when finished. The cloud also supports collaboration and remote workflows, enabling geographically dispersed teams to work as one.

When utilization is below roughly 30%, the economics tilt decisively toward cloud. Beyond cost savings, the real advantages are speed of deployment, experimentation, and innovation. Cloud wins when agility matters most.

On-Demand Capacity Flow



Before – Weeks to Deploy

Capacity planning

Fixed infrastructure



During – Hours to Deploy

Elastic compute


Live events & spikes



After – Nothing to Maintain

Resources shut down

No stranded cost



Control, Predictability, and Always-On Performance

For steady, always-on workloads, on-prem infrastructure remains the gold standard. When channels run 24 hours a day, 7 days a week, the economics of owned hardware — amortized over years — outperform pay-as-you-go cloud rates. Local systems offer consistent latency, deterministic performance, and guaranteed compliance with market-specific regulations.

Leveraging both existing investments and in-house expertise, on-prem environments also provide operational certainty. Engineers can maintain and troubleshoot equipment directly, without depending on external providers or shared-service SLAs. For many broadcasters, this immediate control is mission-critical: regular operations, inbound feeds, and high-value playout chains simply can't tolerate waiting for your cloud vendor to call you back in order to be back on the air.

Owning infrastructure also means predictable budgeting and the confidence that critical services stay online even during a public-cloud outage.


On-prem wins when utilization is constant, performance is non-negotiable, and control is essential.

Availability compared...

99.999% availability is attainable on-prem with an A/B architecture. Public cloud providers generally offer 99.5% to 99.99% SLAs which vary with the individual services offered, making an overall SLA hard to pin down.

To learn more search:

What is the average SLA in {AWS} {GCP} {Azure}



A Unified Control Plane for Every Workflow

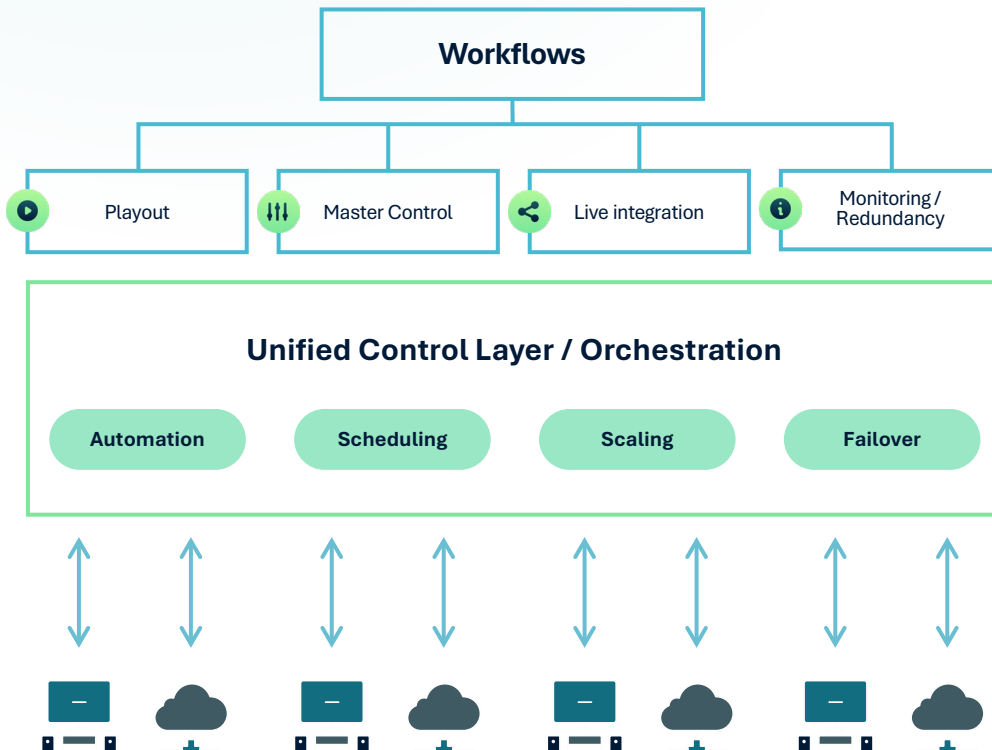
Modern orchestration systems provide a common control layer that unites playout, master control, and live integration across locations and platforms.

Smart orchestration empowers broadcasters to build dynamic, reconfigurable workflows that adapt to changing operational needs. Whether resources are on-premises, in another facility, or in the public cloud, a single interface provides control.

Workflows stay consistent, with shared automation logic and synchronized content, playlists, and metadata, reducing operational risk and eliminating redundant processes. These systems also enable

advanced redundancy strategies and dynamic scaling for events or maintenance windows. Multisite synchronization – linking on-prem and cloud systems together – keeps every channel aligned, with dynamic redundancy ensuring operational continuity even as capacity shifts throughout the day.

Unified orchestration turns multiple infrastructures into one coordinated ecosystem. And orchestration that “abstracts” both cloud and on-prem for the end user, so there is no difference in operation (barring the cloud costs), allows for greater agility, resilience, value, and the lowest TCO.



Where the Hybrid Model Proves Its Value

Across the industry, broadcasters are already putting hybrid principles to work. Some have reduced on-prem load by shifting edit suites and test environments to the cloud, freeing up on-prem resources for mission-critical playout. Others run regional channels in the cloud during prime time and consolidate them overnight to save cost and power.

Hybrid architectures also allow teams to manage channels across multiple sites through a single control layer. For live sports, entertainment, and elections, short-term cloud scaling has replaced the expense and logistics of mobile racks, and overnight shipping. Using hybrid to enable global access and remote collaboration across distributed sites, broadcasters can more effectively consolidate activities and operational overhead across network operations, decrease duplication of work, and reduce overall TCO.

Each of these examples demonstrates that hybrid has become an operational reality.

It provides the flexibility to adapt quickly, optimize continuously, experiment while minimizing risk, and deliver content reliably in any circumstance.




Remote Editing

Shift workflows off site to free local capacity.



Regional Playout

Run regional channels only when needed.



Event Scaling

Spin up channels for major events.



Disaster Recovery

Failover instantly – no duplicate racks.

On-Prem

Control
Performance
Security

Hybrid

Best of Both Worlds

Orchestrated Workloads
Cost Optimization
Operational Resilience

Cloud

Agility
Scale
Remote Access

Hybrid environments combine the agility of the cloud with the control of on-prem infrastructure. When workloads are orchestrated intentionally, organizations can optimize cost, performance, and resilience across both.

Create Operational Resilience Without Expanding the Data Centre

Disaster recovery (DR) is one of the clearest demonstrations of hybrid advantage. Cloud-based redundancy eliminates the need for fully duplicated on-prem infrastructure by allowing broadcasters to keep standby channels and playout engines provisioned but idle, ready to activate in seconds when needed.

If a local failure or planned maintenance event occurs, cloud-based backups can assume control instantly and maintain uninterrupted service. Once normal operations resume, workloads can be scaled back down, avoiding unnecessary cost and energy use.

This flexible DR model not only safeguards revenue and compliance but also provides a continuous testing environment. Systems can be exercised regularly, without disruption, to ensure ongoing reliability.

Where needed, choices for hot, warm, and cold standby system can be made. The highest-value services can have a DR solutions ready for immediate cutover, while lower-value services can be spun up in seconds or minutes. And best of all, orchestration lets broadcasters make these choices on a channel-by-channel basis — and change them at will.

Using cloud resources, it's also possible to rehearse failover processes rather than failing over and hoping things work as intended. This builds confidence with engineering and operational teams, all while keeping spending low.



Hot DR & Redundancy

Channel, content and schedules are continuously synchronized with the main playout channels and always available for immediate failover of parallel running.

Warm DR & Redundancy

Content and schedules are continuously synchronized with the main channels and always available. A channel can be activated in less than a minute from a store configuration.

Cold DR & Redundancy

Content and schedule synchronization is started and begins to upload content and schedules. Full sync may take hours. Channel can be activated in less than a minute from a store configuration though may not be able to play right away depending on content sync status.

Elastic Capacity to Meet Demand

One of the most powerful uses of hybrid architecture is its ability to scale instantly for special events and temporary workloads. Typically, at 5-6% annual utilization, these applications benefit from rapid deployment and global collaboration.

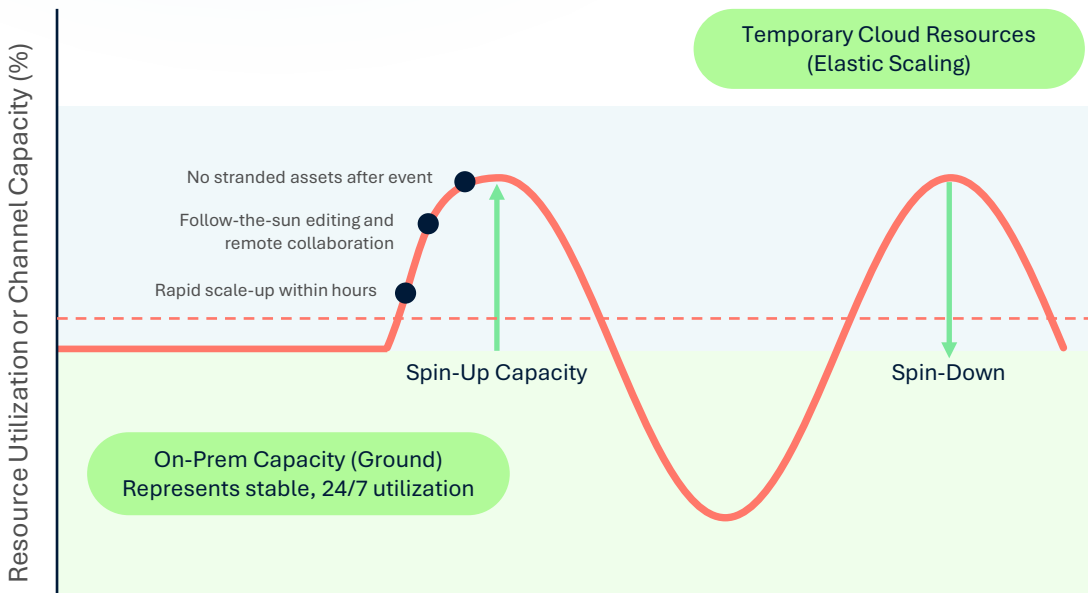
Instead of buying and installing new hardware for a few weeks of programming, broadcasters can use the cloud to spin up additional playout, encoding, or distribution capacity, then release it when the event is over. When utilization drops back to normal, the extra resources disappear, leaving no stranded assets or ongoing costs.

This approach was designed for scenarios including major sporting events, election coverage, and media, where every feed, region, and social stream matters.

Spinning up a few extra channels in the cloud avoids the logistics of shipping gear or deploying staff on-site. It also supports collaboration across time zones, enabling “follow-the-sun” operations, and remote highlight editing.

While hybrid scaling helps broadcasters adapt to unpredictable demand, a hybrid architecture also can help them to better align resources to revenue. Leveraging this approach to support regional splits, broadcasters can sell inventory just to markets where it makes sense.

With the ability to pop up channels only for markets in which they are financially feasible, broadcasters can optimize revenue and cost for revenue.



Time -> Before Event | During Event | After Event

Smart Orchestration, Sustainable Outcomes

With sustainability scores now a part of many RFPs, it's clear that greener offerings—reducing power consumption, reducing e-waste and reducing carbon footprint — have become important across the globe.

Hybrid operations improve efficiency at every level: economic, operational, and environmental. By matching resources to real-time demand, broadcasters reduce both cost and energy waste. Idle servers no longer sit powered and cooled around the clock. Cloud capacity is shared across many users, allowing other organizations use it where there would otherwise be no demand.

This adaptive model extends the useful life of on-prem infrastructure while preventing unnecessary expansion. When workloads move fluidly between environments, every watt and processor cycle serves a purpose. Cloud elasticity also means new services can launch without additional physical infrastructure, helping organizations meet sustainability targets faster. Operationally, broadcasters can establish common workflows that deduplicate effort and boost overall efficiency.

Ultimately, hybrid models enable broadcasters to build resilient systems that consume fewer resources while delivering consistent performance. Along with cost savings, sustainability becomes an outcome of smart engineering and smarter orchestration.



The Path Forward: On-Prem and Cloud in a Unified Operational Model

The future of television is hybrid.

On-Prem and Cloud are not competing technologies; they are partners in optimizing cost and resiliency. Using both, operators can choose the best combination for every job.

Broadcasters already use hybrid workflows every day — for regionalization, live-event scaling, and disaster recovery — and are seeing measurable gains in cost efficiency, energy savings, and operational flexibility.

They leverage elasticity to handle bursty demand, gain diversity to ensure reliability, and are positioned for operational flexibility.

As technology advances, hybrid strategies will advance with it. Smarter orchestration, more intelligent monitoring, and deeper integration between ground and cloud environments will further enhance broadcasters' ability to design for agility, share resources, and optimize overall resource utilization.



The **Hybrid Playbook** for Broadcast

Real-world hybrid implementations show that broadcasters succeed when they place each workload where it performs best — cloud for agility and variable demand, on-prem for predictability and sustained performance — and unify the two through orchestration.

Guided by utilization and designed around resilience, hybrid operations deliver a smarter balance of cost, control, and continuity.

As broadcasters adapt to new pressures and opportunities, hybrid provides the most efficient, resilient, and future-ready foundation for modern media.

Utilization Drives Economics

The crossover point where cloud and on-prem reach cost parity sits around 30% utilization. Below that, cloud's pay-as-you-go flexibility wins; above it, owned infrastructure often prevails.

Orchestration Makes Hybrid Practical

Unified control across ground and cloud ensures consistent workflows, coordinated scaling, and advanced redundancy, turning two infrastructures into one cohesive system.

Hybrid is Resilient by Design

Cloud-based redundancy and event scaling allow broadcasters to respond instantly to outages or surges, reducing downtime and avoiding stranded assets.

Efficiency Equals Sustainability

Right-sized on-prem deployments cut energy and hardware waste, helping organizations meet cost and carbon targets simultaneously.

The Model is Already Here

Broadcasters are using hybrid strategies for regional playout, live events, disaster recovery, and remote operations -- proving that hybrid is the new normal.



Get in Touch

Let's make the business of television better. **Together.**

Visit <http://myimagine.tech/MakeTV>

Sales Inquiries

North America

+1 866 446 2446

Caribbean & Latin America

+52 55 3640 2730 – Mexico

Europe, Middle East & Africa

+44 208 339 1902

Asia Pacific

+852 5819 0900

www.imaginecommunications.com