

HOW AL MASHHAD IS ENGINEERING THE FUTURE OF AI BROADCASTING

In a region where technology and media are converging at speed, Al Mashhad has been quietly building one of the GCC's most ambitious AI programmes – not as a futuristic experiment, but as an integral part of its everyday broadcast ecosystem. In an exclusive interview with Vijaya Cherian, Head of Technology Srinivas Kuppa reveals how the network is redefining what AI means inside a modern newsroom

From its Dubai headquarters, Al Mashhad has engineered what few regional broadcasters have dared to attempt: a fully integrated AI ecosystem that powers the newsroom from ingest to air. Every frame, subtitle and soundbite now passes through intelligent systems designed not to supplant editors and journalists, but to empower them. As Head of Technology Srinivas Kuppa puts it: "The goal has never been to replace human creativity but to amplify it – faster, smarter and always under editorial control."

Launched in early 2024, the initiative has already changed the rhythm of Al Mashhad's newsroom. What began as a set of isolated experiments in video generation and speech recognition has evolved into a hybrid AI infrastructure that handles everything from subtitling and translation to video enhancement and short-form packaging.

"At Al Mashhad, our engineering team adopts a forward-looking approach by viewing challenges as opportunities for innovation. Instead of following the conventional practice

of enumerating popular AI tools, we systematically mapped each challenge to targeted solutions," says Kuppa.

From experiment to enterprise

The company's AI journey formally began in January 2024 with three clear priorities: safeguarding content privacy through local generation, accelerating turnaround from ingest to on-air, and lowering production costs per unit while preserving human oversight. Within months, Kuppa's team moved from local model testing to full-scale deployment.

The first practical milestone had come earlier, however, in 2023, when the creative unit delivered a complex production that combined multiple generative tools and workflows. That experience convinced the network to internalise AI development.


"In 2023, our creative team delivered the Middle East's first AI-driven historical narrative. Building on that momentum, our technology team moved swiftly to develop a proprietary in-house AI tool designed to supercharge future productions

with greater speed, consistency and creative range," says Kuppa.

By Q1 2024, the team had formalised what it calls the Enterprise AI Landscape – a blueprint integrating core production functions with a custom orchestration platform. The roadmap is ambitious: a generative media production suite by late 2024, intelligent content automation integrated with newsroom systems in 2025, predictive operations intelligence embedded directly in broadcast dashboards by 2026.

The multimodal context protocol

At the core of this project is Al Mashhad's proprietary Multimodal Context Protocol (MCP), an in-house platform powered by a custom large language model fine-tuned for Arabic media. It was built to unify every layer of the AI pipeline into one streamlined system, says Kuppa. "The platform was meticulously customised to optimise various broadcast workflow use cases. It operates across multiple GPU servers to enable highly efficient parallel processing,



which is particularly beneficial for task-oriented AI workloads.”

The MCP houses a growing suite of modules: text-to-image, text-to-video, image-to-video, avatar-based video and automatic reels with subtitles; image cropping, highlight generation, upscaling and frame correction; live translation, subtitling and chatbot search; analytics for trending detection and social-video insights. Even the network-operations centre (NOC) now uses a smart contribution downloader that retrieves incoming

“The goal has never been to replace human creativity but to amplify it – faster, smarter and always under editorial control”

SRINIVAS KUPPA, HEAD OF TECHNOLOGY, AL MASHHAD

feeds, automatically enhances them and ingests them into the newsroom system.

Extensive testing of open models on platforms such as Hugging Face enabled the team to fine-tune its own environment for the demands of real-time television. “Al Mashhad’s AI ecosystem spans four domains – content creation, media processing, text-audio-language and video analytics,” says Kuppa.

Hybrid by design

Behind the scenes, Al Mashhad’s infrastructure is deliberately hybrid. Tasks demanding low latency or handling of sensitive content remain on-premises, while scalable or non-critical processes run through cloud APIs.

“Latency and performance for real-time rendering and playout adjacency remain on-premises,” notes Kuppa. “Cost and scalability for batch transcription, translation and classification use local engines scheduled for efficiency. Confidentiality and compliance keep sensitive footage and scripts inside isolated networks with role-based access.”

The stack itself reflects a pragmatic mix of proprietary development and selective partnerships. The engineering team builds on open frameworks such as PyTorch and TensorFlow, layered with custom models optimised for Arabic dialects. Cloud subscriptions handle large-scale media processing, while integrations with ElevenLabs provide text-to-speech in hybrid modes.

“Our AI stack reflects a multi-layered, best-of-breed approach designed for flexibility, scalability and editorial control,” explains Kuppa. “This hybrid ecosystem allows us to balance speed and innovation with security, compliance and oversight”.

Seamless integration

Integrating AI into live production and

AL MASHHAD'S AI INTEGRATION CAPABILITIES

							
Image Creation Generates images from text descriptions	Audio Production Produces audio and speech from text scripts	Video Generation Creates video from text prompts and media	AI Reel Production Produces AI reels with AI-generated subtitles	Content Extraction Extracts video and ideas from social media	Image Cropping Smartly crops images based on interest areas	API Access Provides browser-based ChatGPT API access	Media Upscaling Enhances and upscales AI videos and images

newsroom control was no small feat, requiring workflow re-engineering, staff retraining and new interfaces. Yet by anchoring every deployment in a tangible newsroom need, Al Mashhad minimised disruption. Today, the NOC downloader automates contribution intake, applies enhancement intelligently and pushes feeds into the newsroom system with consistent metadata. Subtitling is now powered by real-time speech-to-text aligned with editorial checkpoints. “Integration was acknowledged as a key challenge, but by anchoring AI deployment in real operational needs – not just tech novelty – we maximised adoption,” Kuppa explains.

Faster, smarter workflow

The impact on daily operations was immediate. AI now accelerates every stage of production. Story research begins with competitive web scrapers that feed prompt generators for ideation; subtitling and captioning are handled by automated engines with human verification; highlight reels and short-form packages are rendered in minutes.

Quantitatively, the gains are striking. Subtitling turnaround has improved by 70%, metadata tagging by 85% and drafting of news summaries by about 40%. “These gains directly support Al Mashhad’s adoption rate for AI-assisted captioning, reels and archival indexing,” remarks Kuppa. “They free journalists

and producers to focus on storytelling, verification and creative development.”

Editors now approve every AI-assisted output, but their roles have evolved from mechanical oversight to creative direction. Producers can generate multilingual versions of stories on demand, while engineers rely on automated QA and predictive monitoring to ensure playout reliability.


Training the human loop

No AI programme succeeds without people. At Al Mashhad, adaptation began with a structured upskilling plan. “AI Tools for Editorial Workflows is a two-

hour hands-on session for producers, editors and journalists,” says Kuppa. “Generative Media and Automation Essentials provides internal knowledge transfer for technical and NOC teams. Responsible AI in News Production raises awareness on ethics and human oversight for all editorial staff.”

Human agency remains a cornerstone of the initiative. “Al Mashhad’s AI principles explicitly prioritise human creativity and mandate human oversight. The mantra ‘machines and humans work together’ reinforces that AI is an assistive tool, not a replacement.”

AL MASHHAD'S ENTERPRISE AI ORCHESTRATION PLATFORM



This cultural positioning has helped neutralise early scepticism. Journalists recognise that automation eliminates drudgery rather than editorial influence, and engineers see that innovation serves operational stability.

For creative teams, AI has become a new sketchpad. Text-to-video and image-to-video tools allow rapid prototyping of visuals; virtual avatars simulate presenters for promos or explainers; and multilingual synthesis enables near-instant localisation.

“Tools like this have expanded creative possibilities,” says Kuppa. “Creative teams now prototype quickly while maintaining brand and tone. Rapid multilingual versions enable earlier feedback from editors and audience teams, improving story framing and reach.”

Guardrails and governance

Technology at this scale brings ethical responsibility. Al Mashhad enforces human review at every step and embeds provenance markers using Adobe Content Credentials-equivalent tags in exported media. “Automation compresses cycle time while mandatory human review, visible attribution and provenance enforcement protect authenticity and journalistic standards,” says Kuppa.

The company’s internal framework defines fairness through access controls and audit trails. Each AI-assisted process includes mandatory checkpoints, while logs ensure transparency. To counter misinformation, the network employs deep-fake detection based on Perplexity analysis and editorial review for all AI-generated or externally sourced material. Copyright and content-ownership issues are handled through strict licensing and storage policies, with clearance logs maintained for audit.

“Planned Milestone 3 includes automated ingest filters running



“By anchoring AI deployment in real operational needs – not just tech novelty – we maximised adoption”

SRINIVAS KUPPA, HEAD OF TECHNOLOGY, AL MASHHAD

source validation, reverse-image checks and metadata consistency scanning,” Kuppa adds.

Compliance and transparency

Ethical guidance for the initiative draws on frameworks from the Dubai Future Foundation Authority, aligning daily practice with transparency and accountability standards. The network has even designed an iconography system – HMC (Human-Machine Collaboration) – displayed across research and generative outputs. “The icon set combines human-machine collaboration symbols with standard research and publishing symbols, so viewers can see at a glance whether an

output was created entirely by humans, created with assistance or generated by machines,” explains Kuppa.

Looking ahead

As the programme matures, Al Mashhad is preparing to extend AI deeper into analytics, recommendation engines and virtual presentation. Upcoming pilots include localised live virtual anchors and synthetic-voice generation for multilingual newscasts. “Near-term priorities include evolving new models integration, AI-driven predictive and proactive operations intelligence, automated video and audio quality control before air, and pilots for localised live virtual presenters,” Kuppa notes.

These ambitions are grounded in measurable editorial value, transparent governance and operational efficiency. In Kuppa’s view, Al Mashhad’s achievement is not merely technical but also philosophical. “This approach demonstrates that a local-first AI programme can deliver measurable editorial and operational value while respecting ethics and transparency. By focusing on security, speed and provenance with human approvals, it offers a practical model that peers in the region can adapt to their own newsroom and playout realities.”

A model for the region

In many ways, Al Mashhad’s initiative encapsulates the next phase of MENA broadcasting: a balance between automation and authenticity. By choosing to build locally, train internally and keep human judgment at the centre, the network has crafted a blueprint that feels both futuristic and familiar, where technology amplifies, not replaces, the craft of journalism.

The result is a newsroom that moves at algorithmic speed without losing its editorial soul. **PRO**