

Copyright Lighting & Sound America

September 2026 complete issue: <https://link.lsamedia.com/sep2026>

When AI Joins the Production Team

By: Dr. Gleb Tsipursky

A production manager asks a generative AI tool to turn scattered notes into a clean load-in checklist. A lighting programmer uses it to summarize fixture documentation. An audio team asks it to organize troubleshooting records from a long tour. These uses can save time, yet they create a basic question for any AI production team: What work can the tool support, and what decisions must remain with a qualified professional?

That distinction matters more in entertainment technology than in many office settings. The practical opportunity sits mostly around the edges of show control. Generative AI can help clean up notes, compare document versions, draft routine communications, classify recurring support questions, and create a first-pass summary of technical material. It can also make mistakes with complete confidence. The right operating principle is simple: use AI to accelerate preparation, while preserving human judgment wherever an output could affect people, equipment, money, contractual obligations, or show continuity.

Start by separating low-risk assistance from high-consequence decisions. Low-risk work includes reformatting a checklist that a person already created, summarizing a manual that the team will verify, or drafting an email that a production manager will review. High-consequence work includes specifying loads, approving rigging configurations, altering safety procedures, interpreting electrical requirements, changing automation behavior, or giving go or no-go direction.

Those decisions belong inside established professional processes, not inside an unverified chatbot response. ESTA's Technical Standards Program exists because industry-specific standards help prevent accidents and protect lives, and an AI tool does not supersede those standards or the qualified people who apply them.

This is where AI guardrails become operational rather than bureaucratic. A useful policy should name approved tools, prohibited information, required review, and clear escalation points. It should also distinguish brainstorming from authorization.

A technician may use AI to suggest possible causes of a network problem, but the technician should verify those possibilities against documentation, measurements, and established troubleshooting practice before changing the system. A production coordinator may ask AI to draft a call sheet, but a responsible person should confirm names, times, access requirements, and safety information before distribution.



Production teams should also label AI-assisted work internally. That label need not become a public disclaimer on every document. It serves a practical purpose: it tells the next person in the chain that the material requires verification.

NIST's framework organizes risk management around governing, mapping, measuring, and managing AI use. Applied backstage, that means defining responsibility, understanding the use case, testing output quality, and correcting failures before the tool reaches a critical production workflow.

The safest place to begin is far from opening night. Test an AI tool on completed productions, archived paperwork, or a rehearsal process where errors carry limited consequences. Compare its output with the work of experienced staff. Track what it gets wrong, what reviewers must correct, and whether it actually saves time after verification. That approach turns AI workflows into measurable experiments rather than acts of faith.

For example, a production company could give an AI system several archived equipment lists and ask it to identify inconsistencies. The exercise would reveal whether the tool catches missing information, invents specifications, overlooks unusual configurations, or produces more review work than it saves.

The company could repeat the test with maintenance notes, venue advances, technical riders, or post-show reports. The goal is to understand where the system performs reliably before anyone depends on it during an active production.

Teams should pay particular attention to handoffs. Live events depend on accurate communication across departments, vendors, venues, performers, and local crews. ESTA's event safety standard for communications emphasizes timely channels for operational, management, security, health, and safety information.

Any AI-assisted message that enters those channels needs technical verification by the person accountable for the underlying information. Speed has value only when the message remains accurate.

Consider an AI-generated load-in summary. It might correctly extract truck arrival times and department contacts while overlooking that a particular entrance has limited clearance. It might shorten a safety instruction until the wording becomes ambiguous. It might combine details from two versions of a document without warning the reader.

The reviewer therefore needs access to the original materials and enough familiarity with the production to recognize what the AI omitted. Asking someone unfamiliar with the project to approve the summary would create the appearance of review without its substance.

Managers also need to make reporting easy. If technicians fear blame for finding an AI error, the organization will learn about failures late. Encourage people to flag bad summaries, invented specifications, missing constraints, and confusing wording. Treat those reports as evidence for improving the process.

This form of human oversight works best when the reviewer has the authority, expertise, and time to reject the output rather than merely rubber-stamp it. A review requirement that people cannot realistically complete offers little protection.

Leaders should resist measuring success by the number of employees using an AI platform. Usage reveals activity, but it says little about whether the system improves the work. Better measures include time saved after review, fewer documentation errors, faster retrieval of technical information, reduced duplication, and the number of AI-generated errors caught before they reached production.

The review burden belongs in the calculation. An AI system that saves 20 minutes of drafting but requires 30 minutes of checking has increased the workload. Teams need evidence from their own workflows instead of accepting broad claims about productivity.

The same logic applies to AI features embedded in production products. *Lighting&Sound America* has already covered a professional audiovisual tracking system that uses computer vision and AI to follow performers in real time.

That kind of production AI can deliver capabilities that a general-purpose chatbot cannot.

Yet embedded intelligence still requires careful commissioning, defined fallback procedures, operator understanding, and testing under realistic conditions. A team needs to know how a system behaves when tracking fails, inputs become unreliable, network connectivity changes, or the operating environment differs from rehearsal.

Every AI-supported process should therefore include a manual fallback. The fallback should remain workable, documented, and familiar to the people expected to use it. A backup process that no one has practiced will offer little help during a live failure.

A good adoption sequence begins with contained administrative uses, expands into supervised technical assistance, and reaches show-critical applications only after disciplined testing. It also gives every tool an owner.

That owner documents the use case, sets review rules, monitors failures, controls significant changes, and decides when the team should stop using the tool. Mature live event safety already depends on defined responsibilities, risk assessment, communication, and standards. AI should enter through that same door.

The larger goal is not to keep AI away from production. It is to prevent convenience from outrunning competence. Teams can capture genuine value when they let AI handle tedious organization while experienced professionals retain control of interpretation and action.

That balance also reduces resistance because employees see a tool designed to support their craft rather than bypass it. Sustainable AI adoption at work grows from that combination of usefulness, transparency, verification, and professional authority.

Before the next production adopts an AI tool, ask four questions. What exact task will it perform? What could go wrong if its output is wrong? Who must verify the result? What happens when the tool fails under pressure?

A team that can answer those questions has the beginnings of a responsible system. A team that cannot answer them has a demo, not a production process. 

—Adapted from: *The Psychology of AI Adoption at Work: From Resistance to Results* (Georgetown University Press, 2026).

Dr. Gleb Tsipursky, a behavioral scientist called the “Office Whisperer” by The New York Times, helps tech-forward leaders stop overpaying for AI while boosting engagement and innovation. He serves as the CEO of the AI consultancy Disaster Avoidance Experts, and has written eight books, including The Psychology of AI Adoption at Work: From Resistance to Results (Georgetown University Press, 2026).