

Study Guide

Module: Create Guardrails for the AI Workflow

This is your working study guide. As you’re going through the course, you’ll have access to one each week. Feel free to save it on your personal drive, take notes and use it during *and* after the course—it will be a great resource now and into the future!

Table of Contents	
Lesson 1: Introduction to AI Governance and Change Management	2
Lesson 2: Establishing Boundaries for Ethical AI Use	9
Lesson 3: Building a Governance Framework	15
Lesson 4: Driving Adoption and Measuring Impact	22
Key Terms	28

In this module, you will:

- Use change management frameworks to communicate the value of automation, align stakeholders, and train teams.
- Establish guardrails for ethical AI use, including oversight, compliance, and brand tone considerations.
- Develop a governance framework and brief change management plan for AI-powered automation.

Lesson 1: Introduction to AI Governance and Change Management

AI-powered automation can drive major efficiency gains—but without the right guardrails, it can also lead to risk, confusion, or wasted effort. In this module, you'll learn how to create responsible, effective workflows by applying the right policies, oversight, and change management strategies. You'll explore key concepts like bias mitigation, human-in-the-loop design, stakeholder engagement, and governance structures that keep your automation aligned with business goals.

To help guide your learning, we'll use the AI workflow guardrails framework—five focus areas that support smart, sustainable automation: purpose, boundaries, roles, monitoring, and support.

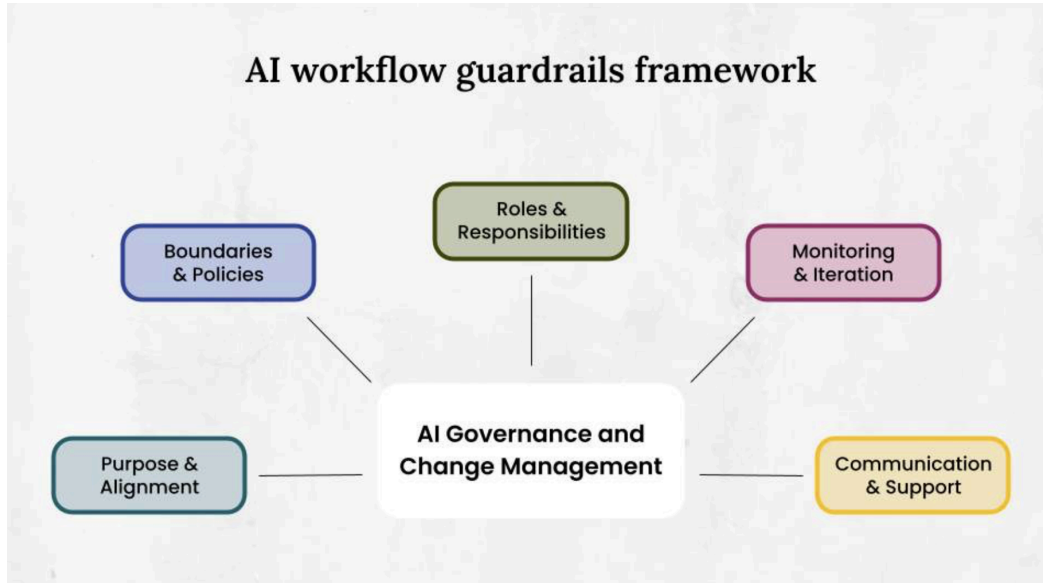


Image description: A diagram titled "AI workflow guardrails framework". This diagram highlights the key components needed to guide and manage AI workflows responsibly through structured governance: "Purpose & Alignment," "Boundaries & Policies," "Roles & Responsibilities," "Monitoring & Iteration," and "Communication & Support."

Notes: What do you hope to learn from this module? How will it help you?

What is AI governance?

A well-designed AI system doesn't just work; it works responsibly. Governance sets the rules and defines what's acceptable, preventing AI from drifting off course.

Defining AI governance

AI governance isn't just a policy; it creates the structure that keeps AI systems aligned with ethical, legal, and business standards. Review the following table to see how governance creates structure and accountability.

Governance component	Definition	Real-world example
Ethical standards	Ensures fair treatment and prevents bias	AI in hiring systems ensuring diversity in candidate selection
Compliance	Aligns AI with legal and contractual requirements	Data privacy policies that protect sensitive customer data
Oversight	Provides monitoring and accountability	Reviewing AI-generated schedules to prevent overwork and burnout

Make the connection: Governance supports the purpose and alignment layers of the AI workflow guardrails framework. It ensures your AI operates with clear intent, stays compliant and ethical, and consistently delivers business value, so automation decisions reflect your goals, not just your data.

What is change management?

Rolling out AI automation is more than a tech shift—it affects people, processes, and decisions. Change management is a critical component of AI success. It keeps teams engaged, informed, and aligned, making adoption smoother and more effective.

Guiding people through change

Even the best AI workflow will fail without team buy-in. Change management helps people understand, support, and adopt new working methods so your AI solution gets used. Review the following benefits of change management:

- **Builds understanding and alignment:**
Helps teams see why AI is being introduced and how it supports their work.
- **Reduces resistance:**
Addresses fears about job loss or confusion by communicating clearly and early.

- **Ensures adoption:**
Guides teams through training and ongoing support so they feel confident using the new workflow.
- **Drives long-term success:**
Reinforces the change over time so people don't slip back into old habits.

Make the connection: Change management is the support layer of the AI workflow guardrails framework. It keeps people informed, confident, and committed, so your AI workflow succeeds beyond the pilot phase.

Change management framework

To guide a successful transition to AI-powered systems, you need more than good intentions—you need a plan. The ADKAR framework offers a step-by-step approach to help people move through change with clarity and confidence. From building awareness to reinforcing adoption, each step plays a critical role in making AI adoption stick.

The ADKAR framework

Rolling out a new AI workflow requires more than just good intentions—it requires a plan. The ADKAR framework gives you a clear, step-by-step path to help people move through change successfully.

- **Awareness:** Communicate why the change is happening and what problem it's solving. Example: Announce the AI rollout early with a clear message about the "why."
- **Desire:** Build motivation by showing how the AI workflow will help individuals and the team. Example: Connect the new workflow to what matters most to each stakeholder group.
- **Knowledge:** Provide the information and training needed to use the new process confidently. Example: Offer hands-on training, reference guides, or demos to explain the new system.
- **Ability:** Support learners as they apply the change in real-world scenarios. Example: Give people time to practice using the new workflow before fully switching over.
- **Reinforcement:** Celebrate wins, share updates, and continue communication to keep momentum going. Example: Send updates that highlight successes and keep the momentum going.

The power duo: Governance and change management

AI governance defines the rules. Change management brings people along. Together, they create the structure that keeps AI aligned with business goals and user needs.

How governance and change management work together

It’s not enough to have policies—or people—alone. For AI systems to succeed, governance and change management must work together. The following table shows how each plays a unique role and collaborates to create lasting, responsible workflows.

Governance focus	Change management focus	How they work together
Sets the rules for ethical AI use	Helps people understand and accept those rules	Ensures adoption of AI that is both safe and trusted
Defines who is responsible for oversight	Prepares stakeholders to take on those roles	Creates clear accountability across the system
Aligns AI with legal/business goals	Aligns people with those goals through communication	Keeps the entire workflow pointed in the right direction

Introducing Soundwave

Imagine you are on the operations team at **Soundwave Entertainment**, a global powerhouse in music and live event production known for bringing unforgettable experiences to fans worldwide. From artist management and tour planning to massive multi-day festivals, Soundwave’s 5,000 employees keep the show running—from stage crews and sound engineers to marketing teams and artist managers.


Soundwave is facing mounting pressure to keep up with the fast pace of live event production. Managing crew schedules, coordinating artist bookings, and aligning venue availability are all handled manually, leading to costly scheduling conflicts, overworked staff, and missed opportunities.

Soundwave is exploring how AI-powered workforce planning can help predict staffing needs, streamline scheduling, and prevent burnout as the company scales. Throughout

this module, we'll use Soundwave as a case study to see how building strong governance and change management practices can lay the groundwork for successful AI adoption.

Keep track of your work

Use this space to take notes and keep track of your work.



Lesson 2: Establishing Boundaries for Ethical AI Use

In Lesson 1, you learned about the AI workflow guardrails framework as a guide for building responsible automation. In this lesson, you'll focus on the boundaries and monitoring layers—how fairness checks, privacy protections, and values alignment help reduce risk and build trust. Keep the framework in mind as you explore how to design AI workflows that are effective, ethical, and clear.

Bias and fairness in AI

Bias in AI can reinforce existing unfairness. The main types of bias—data, algorithmic, interaction, and confirmation—can affect who gets opportunities and who doesn't. It's important to recognize these biases and take steps to prevent them from influencing your AI systems.

Types of bias to watch for

Bias can creep into an AI system in multiple ways. Review the following four common types of bias and what they look like in real-world workflows.

- **Data bias:** Ex: AI reflects historical scheduling patterns that favored certain team members.
- **Algorithmic bias:** Ex: The system prioritizes seniority over specialized skills without justification.
- **Interaction bias:** Ex: Managers override schedules in ways that reinforce old habits.
- **Confirmation bias:** Ex: AI reinforces existing beliefs by filtering out alternative patterns.

How to address bias

Bias in AI is a real risk, but it can be addressed with the proper practices. By building fairness checks into your workflow, you strengthen both accountability and outcomes. Review the best practices below to make fairness a consistent part of your governance strategy.

- Test how your AI impacts different groups across roles and demographics.
- Run regular fairness audits to track shifts, assignments, or access to opportunities.
- Collect user feedback to catch patterns that metrics might miss.

- Assign clear responsibility for monitoring and documentation.

Privacy and compliance in AI workflows

If bias is the first risk, data privacy is the second. Handling personal and company data responsibly builds trust and keeps AI systems compliant with evolving regulations.

Key privacy and compliance practices

AI systems often handle sensitive data, from employee performance to internal documents. These practices help ensure your system is not just legally compliant but also ethical and transparent.

Practice	Example
Minimize use of PII	Only collect data like work history or contact info if it's truly necessary for the task.
Remove sensitive data for training	Strip out client names and financials before using internal documents to train an AI tool.
Ensure transparency and control	Let users see what data is used in scheduling decisions and update their availability or skills.
Follow evolving regulatory standards	Keep up with GDPR, labor laws, or industry-specific rules for how AI can influence employment decisions.
Treat compliance as the floor, not the ceiling	Go beyond the legal minimum to build sustainable, people-centered AI systems.

Transparency is key

Transparency is essential to AI adoption. It builds trust, supports awareness and knowledge, and helps people understand how and why AI makes decisions. In fact transparency and explainability power ethical AI and smooth the path for successful change.

Notes: You just learned that transparency and explainability power ethical AI and smooth the path for successful change. How does that relate to you in your role at work?

Aligning AI with company values

Bias, privacy, and transparency are more than technical requirements—they're core to a company's values. Aligning AI workflows with your organization's mission ensures that every AI decision reflects your brand's values.

Make the connection: Values alignment supports the purpose and support layers of the AI workflow guardrails framework. When AI reflects your organization's core values, it builds trust, strengthens culture, and motivates adoption, turning technology into a meaningful extension of your brand.

Aligning AI with your company's values

Your AI represents your brand, so its decisions should reflect what your organization stands for. These steps help you connect high-level values like fairness or transparency to everyday AI behavior.

Practice: Imagine you're part of the People Ops team at CareBridge Health, a healthcare company known for its commitment to equity, transparency, and employee well-being. You're implementing an AI-powered scheduling system for nursing staff, and your goal is to ensure the system reflects the company's values at every step. Let's explore how they align those values with their AI.

☐ Step 1 - Identify guiding values

Start by selecting core company values that should shape AI behavior, like fairness or transparency.

CareBridge Health example: The team selects fairness, transparency, and well-being as values that guide AI's work.

☐ **Step 2 – Translate values into behavior**

Define what each value means in context.

CareBridge Health example:

- Fairness means equal access to desirable shifts
- Transparency means the AI must explain how shifts are assigned
- Well-being means reducing back-to-back night shifts

☐ **Step 3 – Bake values into design**

Include values from the beginning. Build test cases that ask: “Does this output reflect our values?”

CareBridge Health example: Test cases are created, such as: “Does this schedule reflect our commitment to fairness across teams?”

☐ **Step 4 – Create value-driven statements**

Write clear expectations.

CareBridge Health example: The team documents: “Our AI will prioritize work-life balance over total efficiency when assigning shifts.”

☐ **Step 5 – Give people a way to speak up**

Make it easy for users to raise concerns when the AI doesn’t align with company values.

CareBridge Health example: The scheduling interface includes a feedback tool for staff to flag patterns that feel misaligned with CareBridge’s values.

Keep track of your work

How can you apply this to your job? Jot down your notes from the exercise above in the space provided.

Soundwave’s fairness challenge

Soundwave’s experience shows how fairness risks can surface early—and how the right governance tools can help address them before damage is done.

What happened	Why it happened	How they responded
AI favored bigger venues and top artists	Historical data gave more weight to high-profile events	Flagged it as a fairness issue during pre-launch testing
Small shows and newer crews were deprioritized	The AI mimicked past resource distribution patterns	Implemented bias detection processes to monitor for imbalances
Scheduling choices lacked clear explanations	No visibility into why specific assignments were made	Added transparency protocols to explain AI reasoning

Run a bias check

AI can sometimes reflect bias without you realizing it. This exercise helps you practice spotting those moments—and correcting them.

Your task:

- 1. Open ChatGPT (or a similar tool).
- 2. Input a prompt relevant to your field. Example: *“Write a job ad for a software engineer.”*
- 3. Review the response. Then, enter a follow-up prompt: *“Analyze the response above for potential bias.”*
- 4. Read the tool’s analysis and evaluate whether it missed anything.

Keep track of your work

How could unchecked AI bias show up in your industry? What risks could it create for teams, customers, or business outcomes? Jot down your thoughts in the space provided.



Lesson 3: Building a Governance Framework

This lesson focuses on the roles, monitoring, and support layers of the AI workflow guardrails framework. You'll see how long after launch, identifying risks, defining responsibilities, and building clear escalation paths help ensure AI systems stay aligned with business goals and human oversight.

Thinking through risk

AI-powered workflows come with risks that can undermine their effectiveness. Conducting a SWOT analysis helps pinpoint potential vulnerabilities and strengths, setting the stage for a governance framework that addresses risks while maximizing impact. This risk assessment process connects directly to the *boundaries* and *monitoring* layers of the AI workflow guardrails framework, ensuring your automation stays safe, ethical, and aligned with business goals.

Conducting a SWOT for AI workflows

A SWOT analysis helps you spot where AI can boost your strengths—and where it might introduce risk.

Strengths	Where can AI amplify what's already working well? At Soundwave, AI made generating schedules quickly and consistently across large, complex tours easier.
Weaknesses	Where might existing gaps be amplified by AI? Soundwave had unclear ownership of their automation. Without oversight, issues went unaddressed.
Opportunities	What new capabilities could AI unlock? AI allowed Soundwave to simulate and compare scheduling options before making final decisions—something manual systems couldn't do.
Threats	What could go wrong if the system fails or is misused? Early simulations showed the AI was unintentionally favoring larger venues and top performers, raising fairness concerns.

Evaluating risk

Once you’ve identified potential risks, use these guiding questions to evaluate where governance guardrails are most needed.

Risk consideration	Why it matters
Does the workflow handle personally identifiable information (PII)?	PII must be protected to ensure privacy and compliance.
Will AI influence employment or resource decisions?	Decisions like scheduling or promotions have serious ethical and legal impacts.
Could the workflow introduce or amplify bias?	Bias can damage trust, performance, and fairness.
Is the system mission-critical if it fails?	High-impact systems need strong controls to avoid serious disruption.

Soundwave's governance framework

Soundwave’s AI pilot surfaced a hidden problem: some crew members were overloaded, while others were barely scheduled at all. Without clear oversight, tour managers couldn’t intervene, exposing a major governance gap. Learn how Soundwave used their SWOT analysis to uncover the root cause and build a governance framework with clear roles, escalation paths, and human checkpoints to restore fairness and accountability.

Make the connection: In the previous lesson, you saw how Soundwave used early testing to catch fairness issues before launch. This video picks up where that left off, showing how the team responded when new risks surfaced during their pilot. Together, these moments highlight how governance isn’t just about one-time fixes—it’s a continuous process of adjusting and improving.

From pilot issues to governance fixes

Soundwave’s experience shows how governance becomes critical once AI moves from planning into practice. Here’s what they discovered—and how they responded.

Governance issue	What happened	What they did
No clear ownership of oversight	Crew members were over- or under-scheduled, and no one had authority to fix it	Assigned roles with defined responsibilities for AI oversight
No escalation process for AI decisions	Tour managers couldn't flag or override decisions that didn't make sense	Created formal escalation procedures to intervene when needed
No human review before implementation	Flawed schedules were applied without review	Introduced checkpoints for human review before schedules were finalized

Reflect: Take a closer look at the "What they did" column. These are some specific things you might have to do as you iterate with your workflow. What three things might you have to do as you learn and experiment with your workflow? Jot your thoughts down in your study guide.

Roles and responsibilities in AI oversight

A governance framework is only as strong as the people behind it. Soundwave built a shared responsibility model that assigns clear oversight roles to HR, legal, IT, and operations, ensuring that every risk identified in their SWOT analysis is covered.

Key insights from Soundwave’s oversight model

When roles are clear, oversight becomes actionable. Here’s what Soundwave’s experience teaches us about building a strong, collaborative governance structure.

- **Shared ownership prevents gaps.**
Without defined responsibilities, fairness issues were missed. Assigning oversight roles gave every team a clear purpose.
- **Each team adds a critical lens.**
HR focused on equity, legal on compliance, IT on system reliability, and operations on real-world execution.
- **AI stewards close the loop.**
Trained team members flagged issues early and connected the dots between automation and everyday operations.
- **Oversight became part of the workflow.**
With roles in place, governance wasn’t an afterthought—it was integrated into daily decision-making.

Creating the governance framework: Data

Data is the backbone of any AI governance framework. Soundwave established data protocols to track fairness, privacy, and AI performance, connecting their governance plan to real-world metrics that measure effectiveness.

Building blocks of Soundwave’s data governance

Soundwave didn’t just identify risks—they responded with concrete practices. Explore how they addressed fairness, acceptable use, and privacy by putting clear controls in place across each area of their governance framework.

Issue	What Soundwave did
Fairness	Soundwave discovered their AI was unintentionally creating unequal schedules due to biased historical data. To address this, they implemented three key practices: ● They set fairness benchmarks, like balanced weekly

	hours, travel load, and recovery time by role and experience level. • They launched monthly fairness audits, comparing schedules across demographics and tenure. • They introduced a fairness scorecard reviewed quarterly by department heads to track improvement and accountability.
Acceptable use	To prevent the AI from making inappropriate decisions, Soundwave clarified what tasks were suitable for automation and which required human oversight: • They defined clear rules for acceptable use, such as when the AI could auto-schedule and when human input was required. • Sensitive tasks—like giving performance feedback—were flagged as requiring human review to ensure the AI didn’t overstep.
Privacy	With sensitive employee data in play, Soundwave put strong data protections in place from the start: • They applied data minimization, ensuring the AI only accessed the specific data needed for scheduling. • They explicitly excluded personal characteristics like age and gender from the AI’s inputs. • To reduce long-term risk, they set protocols to delete older data automatically after two years.

Creating the governance framework: People

Human oversight adds a critical layer of protection to AI systems. Soundwave implemented human-in-the-loop checkpoints to prevent biased or inaccurate decisions, reinforcing their governance framework with human judgment and accountability.

Soundwave’s 3-layer oversight structure

Soundwave designed its governance with multiple layers of human-in-the-loop review. These steps created clear pathways for escalation and kept the system responsive and accountable.

- **AI stewards at the front line:**

Trained team members flagged concerns and served as the first point of review.

- **Operational managers review edge cases:**

If a decision seemed questionable, it moved to managers who reviewed real-world fit and context.

- **Governance team handles high-risk decisions:**

Complex issues—like labor law concerns or ethical conflicts—were escalated to the governance team for resolution.

Soundwave's continuous improvement loop

Soundwave built a simple but powerful feedback loop. Here's how they made collecting and acting on input easy.

- **Make it easy to give feedback:** After each scheduling cycle, crew members could rate the results with a quick thumbs-up or down.
- **Facilitate open conversations:** Monthly roundtables allowed people to discuss what was working and what wasn't.
- **Close the loop with action:** A "You said it. We did it." report showed how feedback led to real improvements.
- **Recognize feedback champions:** Team members who suggested improvements were publicly acknowledged for their input.

Notes: Have you ever built a feedback loop? What did it entail and how did it work for you and your team?

Brainstorm your governance framework

Practice: Process mapping

In the module playbook, you'll create an **AI Governance & Change Management Plan**. Let's start sketching out the governance side now.

Your task:

Imagine you're responsible for deploying an AI chatbot at your company. How will you ensure its decisions are safe, fair, and reviewed by the right people?

Keep track of your work

Answer the following questions in the space provided:

- Who should review and approve the chatbot’s decisions or outputs?
- How often should those reviews happen?
- What’s one way to collect user feedback and use it to improve the chatbot?

You'll build on this brainstorm in your final plan, so focus on clarity, not perfection.

Lesson 4: Driving Adoption and Measuring Impact

This lesson focuses on the support and monitoring layers of the AI Workflow Guardrails Framework. While earlier lessons emphasized setting boundaries and assigning oversight, this is about sustaining your AI system through communication, training, and continuous measurement. It's the final piece that ensures your workflow doesn't just launch—it lasts.

Create a change management plan

AI systems can only succeed when people understand the purpose and feel part of the change. Mapping out a change management plan helps identify stakeholders, address concerns, and build a communication strategy that turns potential resisters into project champions.

Change starts with communication

Smart AI implementation starts with people. Let's break down how Soundwave built trust and buy-in through stakeholder planning and targeted communication.

1. **Identify all key stakeholders:** In addition to the oversight team, Soundwave included tour managers and crew members to reflect those most affected by the system.
2. **Assess attitudes and concerns:** They explored each group's sentiment toward the AI and surfaced concerns early, before rollout.
3. **Create role-specific communication:** Messages were tailored to each group's needs, from strategic time savings to shift predictability.
4. **Establish communication rhythms:** Weekly updates, dedicated channels, and virtual town halls gave people a voice and kept everyone informed.
5. **Be transparent about limits:** Soundwave was upfront about what AI could and couldn't do, reinforcing human oversight and control.

Training and support

Training isn't just about clicking through screens—it's about creating confidence in a new way of working. Soundwave's training plan went beyond the basics, helping teams understand AI decision-making and reinforcing human-AI collaboration.

Soundwave’s training strategies

Soundwave’s approach to training was personal, practical, and grounded in how people actually learn and grow.

- **Role-specific paths:** Each group learned what they needed: tour managers focused on oversight tools, crew on schedule use and feedback, and IT on troubleshooting.
- **Training mental models:** Modules helped users understand how AI makes decisions—and when to trust or question the output.
- **Certification program:** Soundwave trained “AI scheduling specialists” to support teams and model strong human-AI collaboration.
- **Real-time and informal support:** Dedicated channels, searchable help articles, office hours, and shadowing gave users help when and how they needed it.
- **Emotional support spaces:** Virtual safe spaces allowed users to express concerns and frustrations—acknowledging the human side of change.

Shadow, pilot, scale: Soundwave’s rollout strategy

Rolling out an AI system all at once can create chaos. Soundwave took a phased approach, testing their AI scheduling system in phases before expanding. Recall how they built trust, gathered feedback, and adjusted.

Soundwave’s phased rollout plan

Soundwave avoided a high-risk launch by rolling out its AI system in deliberate phases.

Phase	What Soundwave did
Phase 1: Shadow Phase	● Purpose: Safely test the system without impacting real operations ● Actions: AI-generated schedules behind the scenes while human schedulers continued as normal ● Built-in strategy: Compared AI results to human schedules to identify needed improvements without any consequences

Phase 2: Pilot Phase	• Purpose: Observe real performance in a controlled environment • Actions: Tested AI scheduling on three different tour types—arena, theater, and festival • Built-in strategies: ○ Safety nets - Clear criteria were set for when to pause or pull back if issues emerged ○ Back-up support - Human schedulers remained on standby to intervene if needed
Phase 3: Staggered Rollout	• Purpose: Scale gradually while resolving issues along the way • Actions: Added five tours at a time with a two-week gap between each wave • Built-in strategy: Lessons learned from each wave were applied before expanding further

Monitoring AI performance for impact

A governance framework means nothing if you’re not measuring impact. Soundwave developed practical metrics to track how well their AI system worked, from schedule accuracy to fairness in crew assignments.

Metrics that matter

Monitoring AI performance isn’t about drowning in data—it’s about choosing metrics that show whether the system is working for people. Review the following table to see how Soundwave did exactly that.

Metric	What it showed	Why it mattered
System uptime	Whether the AI was consistently running	Baseline reliability—ensuring the system was “on” and stable
Prediction accuracy	How closely AI recommendations matched real-world outcomes	Helped validate the system's effectiveness

Schedule stability	Frequency of last-minute changes	Less disruption meant more trust from tour managers
Fairness in assignments	Balance of desirable and undesirable shifts	Reinforced equitable treatment across crew
Manager override frequency	How often humans had to change AI-generated schedules	Fewer overrides signaled increasing trust and AI learning

Iterating and improving AI workflows

AI systems aren't "set it and forget it"—they evolve. Soundwave created a feedback loop to capture real-world insights, prioritize improvements, and keep their AI system aligned with business goals. But how did they keep the system adaptable and effective?

Soundwave's continuous improvement strategy

Soundwave built systems to evolve their workflow over time, turning user feedback into lasting improvements. Review how Soundwave used structured feedback and real-world data to keep improving their AI scheduling system:

- Held weekly "quick fix" meetings to solve urgent issues
- Conducted monthly pattern reviews to spot recurring problems
- Asked quarterly: "Is this still serving the business?"
- Prioritized improvements by effort, impact, and risk
- Made data-backed adjustments, like adding venue fatigue factors
- Improved fairness by adjusting setup times for complex shows
- Closed the loop with feedback by notifying users when suggestions were implemented

Bringing it all together

Soundwave is ready to scale its AI scheduling system across all tours. With their change management and governance frameworks in place, they're focused on creating a continuous cycle of feedback and improvement. Here's how they did it—and what you can learn from their approach.

Soundwave did the following:


- Implemented a careful phased approach
- Created a comprehensive performance monitoring plan

- Focused on real-world outcomes which means they can quickly spot issues and make improvements.
- Kept feedback loops in place

Keep track of your work

By embracing iteration and staying responsive to real-world feedback, Soundwave ensures their AI system continues to evolve, improving performance, building trust, and delivering value long after the initial rollout.

How can you apply this in your job or on your team? What might it look like? Use the space provided to jot down your notes and ideas.



Key Terms

Key Term	Definition	Notes
AI governance	AI governance isn't just a policy; it creates the structure that keeps AI systems aligned with ethical, legal, and business standards.	
Change management	Change management helps people understand, support, and adopt new working methods so your AI solution gets used. Review the following benefits of change management: ● Builds understanding and alignment: Helps teams see why AI is being introduced and how it supports their work. ● Reduces resistance: Addresses fears about job loss or confusion by communicating clearly and early. ● Ensures adoption: Guides teams through training and	

	ongoing support so they feel confident using the new workflow. • Drives long-term success: Reinforces the change over time so people don't slip back into old habits.	
ADKAR framework	Rolling out a new AI workflow requires more than just good intentions—it requires a plan. The ADKAR framework gives you a clear, step-by-step path to help people move through change successfully. • Awareness Communicate why the change is happening and what problem it's solving. • Desire Build motivation by showing how the AI workflow will help individuals and the team. • Knowledge Provide the information and training needed to use the new process confidently. • Ability Support learners as they apply the change in real-world scenarios. • Reinforcement Celebrate wins, share updates, and continue communication to keep momentum going.	
Bias	Bias in AI can reinforce existing unfairness. The main types of bias—data, algorithmic, interaction, and confirmation—can affect who gets opportunities and who doesn't. It's important to recognize these biases and take steps to prevent them from influencing your AI systems.	
Data bias	AI reflects historical scheduling patterns that favored certain team members.	

Algorithmic bias	The system prioritizes seniority over specialized skills without justification.	
Interaction bias	Managers override schedules in ways that reinforce old habits.	
Confirmation bias	AI reinforces existing beliefs by filtering out alternative patterns.	
Change management plan	AI systems can only succeed when people understand the purpose and feel part of the change. Mapping out a change management plan helps identify stakeholders, address concerns, and build a communication strategy that turns potential resisters into project champions.	