

# The AI Incident First-Hour Runbook (Free Lite Edition)

A free artifact of AI GOVERNANCE CHAIN™. The WHAT and the WHY are free. The full playbook, post-mortem kit, and recovery support are the paid path.

## Abbreviations and Acronyms

*No prior knowledge is assumed. Every acronym used in this document is defined below.*

**AI:** Artificial Intelligence

## Who this is for

Anyone responsible for an AI system that could fail in production: a wrong answer at scale, a biased output, a data leak, a fabricated claim, or an autonomous action that should not have happened. This is what to do in the first sixty minutes, before the situation worsens.

## Why this exists (the WHY)

Almost every governance resource plans for prevention. Almost none tell you what to do when prevention has already failed. Yet the first hour of an AI incident determines whether it becomes a contained event or a public crisis. Organizations that have rehearsed the first hour recover quietly. Organizations that improvise often make the damage worse. This runbook gives you the first hour.

## The WHAT: the first sixty minutes

Minutes 0 to 10 - Contain. Reduce the blast radius before you investigate. Can you pause the system, roll back to a safe version, switch to human handling, or restrict who is affected? Containment first, diagnosis second. A system actively causing harm should not keep running while you study it.

Minutes 10 to 20 - Assess scope. Answer three questions. What exactly went wrong? Who is affected, and how many? Is it still happening? Write the answers down. The written record started now becomes your post-mortem and, if needed, your evidence of diligence.

Minutes 20 to 30 - Notify the right people. Reach the accountable owner and whoever must be told: leadership, legal, communications, and, for regulated or public-sector contexts, any party with a notification duty. Notification is not panic. It is control. Surprised stakeholders escalate; informed stakeholders support.

Minutes 30 to 45 - Preserve evidence. Capture logs, inputs, outputs, timestamps, and the system state. Do not clean up before you record. You will need to understand the root cause, and in some contexts you

will need to demonstrate what happened and when.

Minutes 45 to 60 - Stabilize and decide. Is the immediate harm contained? Choose the holding posture: keep the system paused, run it with a human in the loop, or restore it with added safeguards. Set the time for the full review. The first hour ends with the bleeding stopped and a plan named.

## The one principle that governs the first hour

Contain before you explain. The instinct to understand the failure before acting is the instinct that lets the failure spread. Stop the harm first. The explanation can wait an hour; the harm cannot.

## Where the free layer ends and the paid path begins

This lite edition gives you the WHAT (the first-hour sequence) and the WHY (why the first hour decides the outcome). The full edition adds the HOW: the complete incident classification model, the root-cause analysis method, the sycophancy-failure runbook, the post-mortem template, the stakeholder communication scripts, and the regulatory notification guidance.

- - For the full HOW: the AI Incident and Recovery Playbook.
- - For the WITH-YOU: The Governance Circle (incident drills and member support).
- - For the DONE-FOR-YOU: advisory engagement (we help you respond and recover).

Keep this runbook where your team can reach it in a hurry. The first hour arrives without warning.

**Your next step:** Get the [AI-ERF™ Implementation Pack, Full Edition](https://www.aigovernancechain.com/storefront.html) in the Execution Library (www.aigovernancechain.com/storefront.html, section: The Flagship Packs). Every instrument in the Execution Library is also included in the [Governance Vault](#), the all-access annual licence.

---

AI GOVERNANCE CHAIN™

AI Governance Framework

(c) 2026 Mathieu K. Gouanou. All rights reserved.

Member, Harvard Business Review Advisory Council.

Document standard for AI GOVERNANCE CHAIN™: human-first AI governance.