

AI Governance, GRC, and Audit: What Every AI Manager Needs to Know

SEPTEMBER 30

03:00 PM CEST



Ferry Haris

CEO of FEHA International
Consulting B.V.



Islam Elzayat

Founder of GRC by Islam

#GlobalLeadingVoices

The 9:07 a.m. email

MONDAY | EXECUTIVE REQUEST

“Launch an AI hiring assistant across the organization in 90 days.”




DAYS TO LAUNCH

The vendor demo looks excellent. The business case promises faster screening. The deadline is already public.

What must we build around the system?

What would prove it is ready?

The proposed hiring assistant



Five questions become urgent

Purpose

What may the system decide or recommend?

People

Who can benefit or suffer harm?

Ownership

Who can approve, pause, or retire it?

Evidence

What would demonstrate safe operation?

Change

What happens when the model or context changes?

Would you approve the launch today?

Answer in the chat before we show you our decision.

APPROVE

The business value and human reviewer provide enough confidence.

PAUSE

The potential impact requires more work before deployment.

NEED EVIDENCE

The answer depends on controls, testing, and operational proof.


We will return to your vote after the audit.

Three lenses on the same AI system



The AI manager keeps these lenses connected throughout the lifecycle.


The AI manager's accountability



Accountability requires authority to escalate, pause, or retire the system.

ISO/IEC 42001 as the Management system

The standard organizes how the business governs AI. It does not make the business decision for you.



A living management cycle supports daily decisions and certification readiness.

Policy translated into evidence

PRINCIPLES

What the organization commits to

STANDARDS AND RULES

What teams must do

PROCEDURES AND CONTROL POINTS

How work happens

SYSTEM REGISTER AND RECORDS

What the organization knows

OPERATIONAL EVIDENCE

What actually happened


A policy statement has little assurance value until teams can show decisions, tests, approvals, logs, and corrective actions.

Risk, TPRM, and compliance start with context



Standards provide structure. Applicable law, contracts, sector rules, and organizational risk appetite determine the obligations.

Control gates across the AI lifecycle



Every gate needs an owner, acceptance criteria, evidence, and authority to stop progression.

A material change can send the system back to assessment and validation.

The auditor's starting point

A precise audit question keeps the work relevant.

01

OBJECTIVE

What conclusion does the organization need?

02

CRITERIA

Which policy, standard, law, contract, or control applies?

03

SCOPE

Which system, data, vendor, process, location, and period?


04

METHOD

Which interviews, samples, observations, and tests support the conclusion?

Risk-based scope keeps the audit relevant. Evidence-based work keeps the conclusion credible.

The evidence chain



Documents show design. Transactions, behavior, exceptions, and outcomes show operation.

The auditor looks for consistency across all four links.

Signals that governance exists only on paper

- 01  The register omits pilots, embedded features, or employee tools
- 02  The policy names owners who cannot approve or stop systems
- 03  Vendor assurances replace independent validation
- 04  Human review has no time, criteria, training, or escalation
- 05  Dashboards show metrics without thresholds or action

A mature response welcomes these signals because they identify where the system needs improvement.

“Show me the human oversight”

AUDITOR

How does a recruiter challenge the recommendation, and what evidence records that judgment?

CONSULTANT

The procedure requires review. The recruiter can override, and the system records the final choice.

AUDITOR

Good. Now let us sample the logs, decision times, overrides, reasons, and escalations.

The control works on paper, but not yet in practice

ILLUSTRATIVE AUDIT SAMPLE

97%

of sampled recommendations
were accepted within seconds

No recorded rationale. No escalation threshold. No monitoring of automation bias.

CONTROL RESPONSE

- 1 Decision criteria**
Define when recruiters must investigate or escalate
- 2 Evidence**
Capture override and acceptance rationale for selected cases
- 3 Monitoring**
Track acceptance patterns, outcomes, and group impacts
- 4 Authority**
Pause use when thresholds or incidents require action

Conditional pilot after remediation and evidence review

The AI manager's first 30 days

DAYS 1-7

Know the systems

Create or validate the AI inventory. Name the business owner and intended use.

DAYS 8-14

Classify the decisions

Assess impact, applicable obligations, affected people, and risk acceptance.

DAYS 15-21

Connect controls to evidence

Set lifecycle gates. Define required tests, approvals, records, and vendor evidence.

DAYS 22-30

Test the operating reality

Sample one system. Review logs, human behavior, exceptions, and monitoring actions.

Begin with one consequential use case. Build the management rhythm, then scale it.



THANK YOU

✉ ferry@feha.io

✉ islam@horusconsultancy.ca

in <https://www.linkedin.com/in/ferryharis/>

in <https://www.linkedin.com/in/ielzayat/>