

CRAFT the Right AI Harness Around Your Use Case

A practical framework for building an AI harness around a compliance use case



CRAFT QUESTIONS

CONTEXT

WHAT DOES AI NEED TO KNOW?

Do not assume AI knows what your people already know.

WHAT TO CONSIDER

- Policies, procedures, and playbooks
- Decision criteria, thresholds, and exceptions
- Relevant case, request, or transaction information
- Organizational and jurisdictional context
- Trusted data, templates, examples, and prior decisions

PRACTICAL COMPLIANCE EXAMPLES

- G&E policy plus country-specific thresholds and exceptions
- Investigation facts plus relevant policy and report template
- Hotline guidance plus approved response language
- Prior decisions showing how policy is actually applied

REACH

WHAT DOES AI NEED ACCESS TO?

Give AI the access required for the task, not access for its own sake.

- Documents and knowledge repositories
- Compliance, HR, case, or transactional systems
- Employee, third-party, and organizational data
- Workflow, communication, and reporting tools
- Read-only vs. write access and permission limits

- Read the employee's G&E request and relevant policy
- Retrieve information from a case-management system
- Access approved compliance knowledge
- Route a matter into the correct workflow or team

AUTHORITY

WHAT SHOULD AI BE ALLOWED TO DO?

The question is not "Can AI do this?"
It is "Should AI be allowed to do this?"

- Summarize, extract, or draft
- Classify, compare, or analyze
- Recommend or route
- Request additional information
- Act with approval or within defined boundaries

- Draft an investigation summary
- Recommend whether a request meets policy
- Route an employee question
- Approve defined low-risk requests

FAILSAFES

WHEN SHOULD AI STOP OR ESCALATE?

Knowing when to stop is part of being capable.

- Missing or conflicting information
- Low confidence or uncertainty
- Exceptions or novel situations
- High-risk people, transactions, or thresholds
- Matters requiring legal or professional judgment

- Stop when required information is missing
- Route policy exceptions to a compliance professional
- Require human review for high-value or sensitive matters
- Escalate when sources conflict or confidence is low

TESTING

HOW WILL WE KNOW IT IS WORKING?

Trust should be earned through evidence, not assumed at deployment.

- Accuracy and consistency
- Appropriate and missed escalations
- Human overrides and corrections
- Time saved and cycle-time improvement
- Monitoring, feedback, and periodic revalidation

- Compare AI recommendations with expert decisions
- Track human overrides and corrections
- Test known edge cases before deployment
- Revalidate after policy or process changes

Start with the business problem, then **CRAFT** the harness around it.