AI POLICY

The text that says what is done with AI, before the first agent

LEVEL	Organisation
DURATION	1h15 - 2h30
USE	Governance foundation
TRACK	Toolkit

OBJECTIVE

Give the organisation a single document covering two things that generic policies rarely separate: the daily use of AI tools, and the moment a system decides in place of a human. It is signed, it is revised, it comes before the first agent.

1 THE RULE · TWO REGIMES, ONE TEXT

USE AND DECISION

Assisted AI proposes and the human decides. Agentic AI decides and acts within a perimeter. Legal accountability sits in the second regime, which almost no usage policy addresses. The ACF policy covers both and defers to the agentic constitution (ACF-04) for the detail of each agent.

WHAT THE POLICY COMES BEFORE

It is laid down before the first agent and grounds the mandates (ACF-13) and the non-delegable zones. An organisation that writes mandates with no policy writes rules and nobody knows on whose authority they bind.

2 TO BE FILLED IN

PART	WHAT IT SETS	WHO OWNS IT	REVIEW

3 CHECK WHEN RE-READING

THE LIST OF PROHIBITIONS

It must name decisions, not technologies. A ban aimed at a product ages within six months; “no account closed without a human” holds.

THE REVIEW DATE

A policy with no review date written in becomes an archive document the day the tooling changes.

A policy that does not list the decisions it refuses to delegate is not a policy, it is an intention.

© 2026 Vincent Dorange · acfstandard.com/licence