INTELLISYNC SOLUTIONS

STRUCTURE - CLARITY - BETTER DECISIONS

DECISION ARCHITECTURE

Decision Guardrail Canvas

A worksheet for defining what an AI system may recommend, what it may execute, and where human judgment remains mandatory.

WHAT THIS TEMPLATE HELPS YOU DECIDE

Convert responsible AI from broad policy into operating rules: system role, decision rights, escalation thresholds, forbidden actions, evidence requirements, and review cadence.

BEST FOR

- Teams moving from AI experiments into recurring operations
- Customer, sales, finance, HR, or advisory workflows with business risk
- Leaders who need clarity before delegating work to agents or copilots

OUTPUTS

- A clear decision-rights boundary
- Escalation rules for ambiguous or high-risk outputs
- A review model that keeps accountability visible

STEP 1

Define the system role

Every AI-native workflow needs a role boundary. The system should know whether it is drafting, classifying, recommending, routing, or executing.

Role definition

Primary operating role

Example: draft response, summarize account risk, classify inbound requests.

Business outcome supported

Name the operational result this system improves.

Accountable owner

Name the person or role responsible for decision quality.

Allowed operating modes

- [ ] The system may summarize approved context
- [ ] The system may classify or prioritize records
- [ ] The system may draft recommendations for review
- [ ] The system may route work to an owner or queue
- [ ] The system may execute only pre-approved low-risk actions

STEP 2

Separate recommendation from execution

AI governance becomes practical when teams can distinguish decision support from autonomous action. This page defines that boundary.

Decision-rights map

Recommend	What may the AI propose without committing the business?
Draft	What may the AI prepare, but not send, approve, price, or finalize?
Route	Which tasks may the AI assign or escalate based on deterministic rules?
Execute	Which actions, if any, may run automatically after validation?
Never execute	Which actions must remain human-owned regardless of confidence?

Determinism beats inference. If an operation can be handled by a validated tool, workflow rule, permission check, or data contract, do not leave it to model judgment.

STEP 3

Set escalation rules

Escalation rules keep ambiguity from becoming invisible automation. Use them to protect customers, staff, and the business when the system reaches uncertainty.

Escalation triggers

- [ ] The request includes legal, financial, HR, safety, or regulated advice
- [ ] The output changes a customer commitment, price, timeline, or obligation
- [ ] The model confidence, source coverage, or evidence quality is weak
- [ ] The customer sentiment is angry, distressed, urgent, or adversarial
- [ ] The workflow involves sensitive data or cross-functional approval

Review model

Reviewer	Who reviews the output before it can move forward?
Approval proof	What evidence shows the review happened?
Override path	What happens when the reviewer rejects or edits the AI output?
Cadence	How often are exceptions reviewed for pattern detection?