LIBRA SENTINEL

Data Privacy & AI Governance

LLM-Assisted Creativity for Workflows

AI ACT
ARTICLE 4
AI LITERACY

1) Privacy-notice variants (web / app / email)

Idea seed (current version) → LLM remix → Diversity prompt → Provenance tag → Legal review (owner: DPO).

Why it matters: Deliver legally compliant notice copy tuned to different audiences and UX friction levels (legal-safe → plain English → ultra-short banner).

Prompt (copy/paste):
(A) Full legal text for the policy page (include all required Article 13/14 disclosures).
(B) Plain-English summary for in-app screens (≤60 words).
(C) Ultra-short banner for email (≤20 words + one CTA/link to full policy).
(D) Just-in-time pop-up for specific actions (e.g., location/microphone) stating: what is collected, why, and an Allow/Deny choice with a "Learn more" link.
For each variant, list which required disclosures are covered and add one UX note about consent flow/friction.
Outputs to keep: 3 variants + required-disclosures checklist + UX note.

Converge/check: Compare the variants for readability and compliance coverage. Keep the ones that balance low friction + legal sufficiency, and verify disclosures with your DPO/legal team.

2) DPIA brainstorming (scoping & showstoppers)

Idea seed (project brief) → LLM flow map → Anti-convergence regeneration → Showstopper flag → Escalate if flagged.

Why it matters: Rapidly surface purposes, data flows, mitigations, and any DPIA "showstoppers" before a costly implementation.

Prompt (copy/paste):
For [PROJECT], list:
(1) three plausible legal/operational purposes,
(2) an end-to-end data flow (sources → processors → storage → retention),
(3) five mitigation options for privacy/security risks, and
(4) any DPIA showstoppers that would likely require stopping or redesigning the project.

Outputs to keep: purpose list, flow diagram text, prioritized mitigations, and explicit showstoppers.

Converge/check: If any showstopper exists, escalate; otherwise keep top 3 mitigations and assign owners for DPIA drafting.

AI ACT
ARTICLE 4
AI LITERACY

3) Vendor-intake triage (scoring & deal-breaker questions)

Idea seed (vendor answers) → LLM rubric table → Provenance + scoring → If score < threshold or deal-breaker = escalate.

Why it matters: Turn onboarding forms into instant triage tools that flag high-risk suppliers and standardize decisioning.

Prompt (copy/paste):
Create a vendor-triage rubric for [VENDOR TYPE]: include
(A) 6 scoring criteria (divergent options) (e.g., data residency, SOC/ISO attestation, subcontractors, breach history, contractual indemnities, privacy controls) with 1–5 scoring rules, and
(B) 5 hard deal-breaker questions that trigger escalation.

Outputs to keep: rubric table (criteria + scoring), threshold for escalation, and deal-breaker list.

Converge/check: If vendor scores below threshold OR any deal-breaker = escalate to procurement & legal; otherwise produce a tailored remediation plan.

4) Incident tabletop (diverse scenarios + comms trees)

Idea seed (system) → LLM scenarios → Compare human & AI scenarios → Keep top human + top AI + owner for playbook.

Why it matters: Prepares teams for varied breach shapes and gives immediate notification and press-line templates.

Prompt (copy/paste):
Generate 4 diverse breach scenarios for [SYSTEM]:
(A) data-exfiltration via third-party,
(B) misconfiguration exposing PII,
(C) insider misuse,
(D) AI model leak.
For each: timeline of events, likely impacted stakeholders, primary containment steps, notification tree (who to call in order), and one draft press line (≤30 words).

Outputs to keep: scenario cards, notification trees, and draft press lines.

Converge/check: Choose the top scenario(s) matching current controls gaps and add the containment owner & 24-hr checklist to the incident runbook.

AI ACT
ARTICLE 4
AI LITERACY

5) AI-policy options (BYOD + Gen-AI usage models)

Idea seed (current policy) → LLM drafts (3 postures) → Diversity prompts (restrictive→permissive) → Choose & assign enforcement owner + KPI.

Why it matters: Generates alternative policy models (restrictive → conditional → permissive + monitoring) so leadership can pick a governance posture and enforcement model.

Prompt (copy/paste):

Draft 3 AI-use policy options for employees (BYOD + generative AI):
(A) Restrictive: only approved tools,
(B) Conditional: allowed with controls (data redaction, no client PII),
(C) Permissive with monitoring: allowed plus logging and audits.
For each option, list enforcement steps, pros/cons, and a short executive summary (1 paragraph).

Outputs to keep: 3 policy drafts, enforcement matrix, and exec summary.

Converge/check: Pick the policy whose pros/cons fit risk appetite; define enforcement owner and 90-day review metric (e.g., % of incidents tied to non-compliant use).