

LPM Reusable Decision Log

A versioned knowledge object for decision memory, authority, evidence, AI boundaries, and governance traceability.

Object type	LPM Knowledge Object
Primary LPM layers	Decision Architecture, Governance Architecture, Information Ecology, AI Amplification
Connected layers	Ownership Map, Communication Architecture, Platform Structure, Identity & Incentives
Primary use	Create a durable decision memory for material decisions, including owner, options, evidence, rationale, AI involvement, downstream impact, review date, and supersession history.
Website use	Downloadable template, executive guide, decision workshop worksheet, JSON object for Lapemo ingestion, and future guided skill.
Version	1.0
Owner	LPM / Lapemo
Last reviewed	2026-06-23

A Decision Log creates organizational memory. It records what was decided, who had the authority to decide, why the decision was made, what evidence was used, what tradeoffs were accepted, how the decision was communicated, and when it must be reviewed.

The object is reusable across three company profiles: 500+, 5,000+, and 10,000+ employees.

1. Knowledge object model

What this knowledge object contains

- Human guide - plain-language guidance for executives, transformation leaders, product, technology, risk, operations, data, and AI leaders.
- Template artifact - a downloadable decision log teams can use during strategy, planning, governance, delivery, and AI scaling.
- Machine-readable schema - fields, scoring logic, validation rules, mappings, prompts, review triggers, and version metadata.
- Guided skill - an AI-assisted workflow that classifies decisions, detects missing accountability, creates decision summaries, and flags stale or conflicting decisions.

Why this matters for AI scaling

AI makes weak decision memory dangerous. When decisions are scattered across meetings, chats, documents, dashboards, and system tickets, AI cannot safely reason over enterprise intent. A decision log becomes the source of truth for what the organization actually chose and why.

2. Canonical fields and scoring

Required fields

Field	What it captures	Required?
Decision ID	Unique identifier for the decision record	Yes
Decision title	Short human-readable name for the decision	Yes
Company profile	500+, 5,000+, or 10,000+ employee version	Yes
Decision type	Strategic, operating, technical, risk, funding, product, people, data, platform, or AI decision	Yes
LPM layer	Primary LPM layer affected by the decision	Yes
Decision owner	One accountable owner for the decision record and outcome	Yes
Decider	Role or person with authority to make the decision	Yes
Recommenders	Roles or teams that shaped the recommendation	Yes
Consulted	Roles that must be consulted before finalization	Required when cross-functional
Informed	Roles or audiences that must be informed after decision	Yes
Decision date	Date the decision was made	Yes
Effective date	Date the decision becomes active	Yes
Status	Proposed, approved, rejected, deferred, active, superseded, retired	Yes
Decision context	Problem, opportunity, or constraint that triggered the decision	Yes
Options considered	Alternatives evaluated before the decision	Yes
Chosen option	The selected path	Yes
Rationale	Why this option was selected	Yes

Field	What it captures	Required?
Evidence inputs	Data, analysis, policy, customer signal, risk input, or AI output used	Yes
Assumptions	Known assumptions behind the decision	Yes
Tradeoffs	What the organization knowingly gave up	Yes
Risks accepted	Risks explicitly accepted by the decider	Required when risk exists
Downstream impacts	Teams, workflows, systems, customers, employees, or controls affected	Yes
Systems impacted	Platforms, records, tools, models, or workflows touched	Required when systems change
AI involvement	None, assisted, recommended, drafted, routed, executed, or autonomous	Yes
AI decision boundary	What AI can and cannot do in relation to the decision	Required when AI involved
Human review owner	Owner accountable for human review and override	Required when AI involved
Control owner	Risk, compliance, security, legal, privacy, model risk, or audit owner	Required for material or regulated decisions
Communication path	How the decision is communicated and where it is stored	Yes
Implementation owner	Owner accountable for execution	Yes
Value metric	Metric that proves the decision worked or did not work	Yes
Review date	Date to revisit or validate the decision	Yes
Revisit trigger	Condition that forces review before the review date	Yes
Supersedes decision ID	Prior decision replaced by this decision, if applicable	Optional
Source links	Links to artifacts, meeting notes, Jira, docs, approvals, dashboards, or evidence	Yes
Version	Object version and effective date	Yes

Decision log scoring logic

Score	Meaning
0	Decision is not logged, owner is missing, or rationale is unavailable
1	Decision is captured, but owner, rationale, evidence, or communication path is incomplete
2	Decision has owner, rationale, and date, but options, evidence, impact, or review rules are weak
3	Decision has clear authority, evidence, rationale, impacts, implementation owner, and review date
4	Decision is versioned, linked to evidence, monitored, communicated, supersession-ready, and AI-safe

Decision memory score: average of ownership clarity, decision authority, evidence quality, rationale quality, impact mapping, communication completeness, review discipline, and AI boundary clarity.

High-risk flag: any material decision with a score below 3, AI involvement above Level 1, or missing control owner when risk is medium, high, regulated, or critical.

3. Role model

Role	Accountability	Approval responsibility
Decision owner	Owns the decision record, outcome traceability, and review discipline	Yes

Decider	Has the authority to make or approve the decision	Yes
Recommender	Creates the recommendation, options, evidence, and tradeoff analysis	No
Consulted role	Provides input that must be considered before the decision	No
Informed role	Receives the decision after approval	No
Implementation owner	Turns the decision into execution, system change, workflow change, or policy change	Yes
Control owner	Confirms risk, legal, compliance, security, privacy, model risk, or audit implications	Required for material decisions
Human review owner	Owns review, override, escalation, and exception handling when AI is involved	Required when AI is involved
Value owner	Tracks whether the decision delivered the intended outcome	Recommended

Owner rule

Every material decision needs one accountable decision owner. Committees can advise, review, or approve. They cannot replace accountable ownership.

4. Version A - 500+ employee company

A 500+ employee company usually moves fast, but decision memory often lives in leaders, Slack, Teams, and recurring meetings. The log should be lightweight enough to use, but strong enough to prevent ambiguity.

Design area	500+ version	Minimum standard
Decision scope	Team, product, function, or leadership decision	Keep the log simple and visible
Owner model	One accountable decision owner and one implementation owner	No decision without a named owner
Evidence	Customer signal, operator judgment, financial data, delivery impact, risk note, or leadership input	Do not require heavy bureaucracy for low-risk calls
Communication	Post decision in one visible location and inform affected teams	Reduce Slack and meeting memory dependence
AI involvement	Capture whether AI drafted, summarized, recommended, or automated part of the decision	Any AI-assisted decision needs human owner
Review cadence	Monthly or quarterly for active decisions	Prevent stale decisions from becoming tribal rules

500+ design principle: keep the decision log simple, visible, and owner-led. The goal is not bureaucracy. The goal is to stop decisions from disappearing into tribal memory.

5. 500+ implementation notes

Recommended operating pattern

- Use one shared decision log for leadership, product, operations, technology, and AI initiatives.
- Require a decision owner, rationale, communication path, and review date for any material decision.
- Log AI-assisted decisions even when AI only drafted, summarized, or recommended.
- Review active decisions monthly during operating reviews.
- Mark superseded decisions clearly instead of deleting history.

Common failure modes

Failure mode	Signal	Fix
Decision lives in a meeting	Teams ask the same question repeatedly	Create a short decision record with owner, rationale, and source links
Founder or executive memory becomes the system	People wait for verbal clarification	Move decisions into a visible log
AI recommendations become undocumented influence	Teams cannot explain why a path was chosen	Capture AI involvement and human decision owner
No review date	Old decisions keep driving work after assumptions changed	Add revisit trigger and review cadence

6. Version B - 5,000+ employee company

A 5,000+ employee company needs decision memory across functions. The log should connect business direction, technology execution, risk input, data evidence, funding, product priorities, and platform dependencies.

Design area	5,000+ version	Minimum standard
Decision scope	Cross-functional, product, technology, risk, data, funding, operating, or platform decision	Use one shared log across major initiatives
Owner model	Decider, decision owner, implementation owner, control owner where needed	Separate decision authority from execution ownership
Evidence	Roadmap, business case, customer data, risk analysis, architecture input, data quality, or policy source	Require evidence for material decisions
Communication	Decision summary linked to Jira, Confluence, governance notes, and leadership forums	Make decisions searchable and reusable
AI involvement	Classify AI use and decision boundary	AI recommendations require human decision owner and evidence trail
Review cadence	Quarterly portfolio review plus event-driven review	Review when assumptions, systems, risk, or incentives change

5,000+ design principle: make the decision log the shared memory layer between business strategy, product delivery, technology, data, risk, and governance.

7. 5,000+ implementation notes

Recommended operating pattern

- Connect major decision records to roadmap epics, governance forums, architecture decisions, risk reviews, and operating metrics.
- Require options, rationale, evidence, decision owner, implementation owner, and affected systems for material decisions.
- Use the decision log in quarterly business reviews and portfolio reviews.
- Require AI decision boundary classification for AI-assisted or AI-enabled decisions.
- Use supersession history to reduce conflicting direction across departments.

Common failure modes

Failure mode	Signal	Fix
Departments make conflicting decisions	Teams execute against different assumptions	Create one shared decision ID and supersession model
Evidence is disconnected from decision	Leaders debate the same facts repeatedly	Link evidence, assumptions, and rationale directly to the record
Risk is consulted too late	Controls appear after delivery has momentum	Add control owner and risk review before active status
Decision and execution split	Nobody owns follow-through	Add implementation owner and value metric

8. Version C - 10,000+ employee company

A 10,000+ employee company needs an enterprise decision ledger. Local teams need speed, but the enterprise needs auditability, legal and regional awareness, risk traceability, AI boundary control, and decision lineage.

Design area	10,000+ version	Minimum standard
Decision scope	Enterprise, BU, region, legal entity, platform, risk, data, AI, or customer-impacting decision	Federate local decisions into an enterprise ledger
Owner model	Decider, accountable decision owner, implementation owner, control owner, value owner, and jurisdiction owner where needed	Authority must be auditable by role and level
Evidence	Formal evidence pack, control evidence, audit trail, data lineage, model or agent trace, and approval history	Material decisions need replayable evidence
Communication	Decision ledger connected to governance, architecture, risk, roadmap, model registry, data catalog, and operating updates	No material decision should live only in a meeting deck
AI involvement	Document decision boundary, human review, model or agent owner, monitoring owner, and prohibited actions	Autonomous or regulated AI requires formal approval
Review cadence	Tiered review by risk, region, system, model, and decision boundary	Supersede or retire stale decisions explicitly

10,000+ design principle: federate decision-making, but centralize decision memory, evidence standards, AI boundaries, and supersession history.

9. 10,000+ implementation notes

Recommended operating pattern

- Maintain a federated decision ledger by business unit, function, region, legal entity, platform, initiative, and risk tier.
- Connect decision records to policy, controls, model registry, architecture repository, data catalog, roadmap, and audit evidence.
- Require formal evidence packs for high-impact, regulated, customer-facing, employee-impacting, financial, or AI-autonomous decisions.
- Create decision conflict detection across related initiatives and platforms.
- Use explicit supersession and retirement rules so old decisions do not remain operationally active.

Common failure modes

Failure mode	Signal	Fix
Local decisions become enterprise conflicts	Regions or BUs implement incompatible direction	Use federated ledger with global decision IDs and local context
Audit cannot reconstruct why a decision happened	Evidence sits across decks, emails, chats, and tickets	Create evidence pack and version history
AI agents act on stale decisions	Automation follows old policy or outdated assumptions	Require current active decision source and review trigger
Old decisions never retire	Teams cite outdated approvals	Use superseded and retired statuses with lineage

10. Decision log template

Use this table as the editable working template. In Lapemo, these same fields become structured records with validation rules, evidence links, ownership mappings, and review triggers.

Field	Example value	Instruction
Decision ID	DL-0001	Unique and stable ID
Decision title		Short title
Decision type		Strategic, operating, technical, risk, funding, product, people, data, platform, AI
Decision owner		One accountable owner
Decider		Person or role with authority
Status		Proposed, approved, active, superseded, retired
Decision date		YYYY-MM-DD
Effective date		YYYY-MM-DD
Context		Why this decision is needed
Options considered		At least two options for material decisions
Chosen option		Selected path
Rationale		Why this path won

Field	Example value	Instruction
Evidence inputs		Links or references
Tradeoffs		What was sacrificed
Risks accepted		Explicit risk acceptance
Downstream impacts		Teams, workflows, systems, controls, customers, employees
AI involvement		None, assisted, recommended, routed, executed, autonomous
AI decision boundary		What AI can and cannot do
Communication path		Where the decision is stored and announced
Implementation owner		Who makes it real
Value metric		How success will be judged
Review date		When it must be revisited
Revisit trigger		Condition that forces review

11. AI decision boundary levels

Level	AI role	Required control	Human accountability
Level 0	No AI used	Normal decision log only	None
Level 1	AI summarizes or drafts context	Decision owner validates accuracy	Human accountable for final wording
Level 2	AI recommends options	Evidence and assumptions must be visible	Human decider chooses and records rationale
Level 3	AI routes or triggers workflow steps	Process owner and control owner approve flow	Exceptions and override path required
Level 4	AI executes with human approval	Control evidence and monitoring required	Human approval required before action
Level 5	AI acts autonomously	Formal governance approval required	Autonomy, monitoring, incident path, and stop control required

12. Validation rules

Validation condition	System response	Reason
Missing decision owner	Block approval	Every material decision needs one accountable owner
Missing rationale	Block active status	A decision without rationale becomes politics or memory
AI involved but no boundary	Block production or execution	AI scope must be explicit
Material risk but no control owner	Escalate	Risk acceptance must have accountable control owner
No review date	Flag stale risk	Decisions age and assumptions drift
Superseded decision still active	Require cleanup	Old decisions must not compete with new decisions
No communication path	Flag adoption risk	A decision is not real until affected teams can find it

13. Mapping rules

Connected LPM object	Decision log mapping
Ownership Map	Decision owner, decider, implementation owner, control owner
Decision Rights Matrix	Decision type, authority level, consulted roles, informed roles
Incentive Alignment Checklist	Value metric, behavior change, tradeoffs, risk acceptance
AI Initiative Owner Register	AI involvement, AI boundary, owner, risk tier, human review owner
Governance Architecture	Approval status, control owner, evidence, review date, escalation path
Information Ecology	Evidence inputs, source links, assumptions, decision memory
Platform Structure	Systems impacted, workflow changes, data sources, implementation dependency
Communication Architecture	Communication path, informed audiences, decision summary

14. AI prompts for the guided skill

- Classify this decision by type, LPM layer, risk level, and company profile.
- Identify the decider, decision owner, implementation owner, consulted roles, informed roles, and control owner if needed.
- Summarize the context, options considered, chosen option, rationale, evidence, assumptions, and tradeoffs in plain language.
- Detect whether AI influenced the decision and classify the AI decision boundary from Level 0 to Level 5.
- Find missing fields that would prevent this decision from being approved, implemented, or audited.
- Generate an executive decision summary, a team communication note, and a governance-ready record from the same object.
- Detect stale, superseded, conflicting, or duplicated decisions across an initiative, function, or enterprise portfolio.

15. Render outputs

Output	Use
Word document	Workshops, consulting, internal enablement, website download
PDF guide	Executive briefing, education, static reference
Website page	Public resource hub and SEO content
Interactive form	Guided decision capture and validation
JSON object	Lapemo ingestion and versioned knowledge object storage
CSV or table view	Bulk import into decision ledger, Jira, Airtable, Notion, or governance tools
In-app workflow	Live decision memory, stale decision alerts, conflict detection, AI boundary governance

16. Review and update triggers

- Decision reaches review date.
- Decision owner changes role or leaves the organization.
- AI begins influencing, recommending, routing, or executing work related to the decision.
- A downstream system, data source, policy, workflow, model, agent, or vendor changes.
- A risk, control, legal, compliance, privacy, or security issue appears.
- A newer decision supersedes the existing decision.
- Measured outcomes show the decision is not producing intended value.

17. Website positioning copy

Decision Log

A reusable LPM template for turning decisions into durable organizational memory. Use it to capture ownership, authority, rationale, evidence, tradeoffs, downstream impact, AI involvement, and review triggers before decisions become scattered across meetings, chats, decks, and systems.

18. Practical guidance

A decision log is not meeting minutes. It is an operating control. It should be short enough that teams actually use it and structured enough that the organization can later understand, audit, automate, and improve decisions.