

LPM REUSABLE ARTIFACT

Knowledge Freshness Review

A living knowledge-object template for reviewing whether artifacts, policies, prompts, dashboards, playbooks, and AI-readable content are still current, trusted, and safe to use.

Field	Value
Object type	LPM Knowledge Object
Primary LPM layers	Information Ecology, Governance Architecture, AI Amplification
Connected layers	Ownership Map, Decision Architecture, Communication Architecture, Platform Structure, Evidence Checklist, Source-of-Truth Map
Primary use	Review whether knowledge objects, artifacts, policies, prompts, dashboards, playbooks, and AI-readable content are still accurate, owned, approved, and safe to use.
Website use	Downloadable template, website resource, AI readiness artifact, JSON object for Lapemo ingestion, and future guided review skill.
Version	1.0
Owner	LPM / Lapemo
Last reviewed	2026-06-24

Purpose

The Knowledge Freshness Review defines how an organization keeps its operating knowledge from becoming stale, duplicated, unsafe, or misleading. In LPM terms, it keeps Information Ecology healthy enough for decisions, governance, platform operations, and AI amplification.

This is not a document cleanup exercise. It is an operating-control practice for companies scaling AI across human and machine workflows.

Core principles

Principle	Meaning
Knowledge expires differently by risk	A framework definition may last a year. An AI retrieval source, policy, pricing rule, system integration, or decision rule may become unsafe in days.
Freshness is operational control	A stale document is not just a content problem. It can create bad decisions, broken handoffs, weak governance, and unsafe AI outputs.
Every knowledge object needs an owner	A page, template, prompt, dashboard, playbook, or policy without an accountable owner will drift as the company changes.
AI should not learn from abandoned artifacts	Any object that is retrievable by AI must carry status, owner, version, freshness window, and human review rules.
Stale does not always mean wrong	Stale means the object has passed its approved confidence window and must be reviewed before it is treated as decision-grade.
Review is a workflow, not a yearly cleanup	Freshness should be triggered by time, system changes, org changes, policy changes, incidents, AI usage, and evidence gaps.
Supersession matters	When a knowledge object changes, old versions should be marked superseded or retired so teams and AI agents do not use competing truth.

Required fields

Field	Definition	Required
Knowledge object ID	Unique identifier for the artifact, page, prompt, policy, template, playbook, metric, or AI-readable source	Yes
Object name	Human-readable name of the knowledge object	Yes
Object type	Template, checklist, policy, SOP, decision rule, prompt, dashboard, metric, data definition, playbook, guide, or website page	Yes
Primary LPM layer	The layer the object governs or supports	Yes
Intended audience	Executives, transformation leaders, operators, HR, IT, governance, data teams, AI owners, or external users	Yes
Business purpose	What decision, workflow, onboarding step, control, or AI use case this object supports	Yes
Accountable owner	Person or role accountable for accuracy, approval, and lifecycle	Yes
Knowledge steward	Person or role responsible for review execution, metadata hygiene, and publication workflow	Required when material
Source-of-truth location	Canonical system or URL where the approved object lives	Yes
Published formats	Website page, DOCX, PDF, Markdown, JSON, form, CSV, in-app workflow, or AI retrieval chunk	Yes
Current status	Current, under review, stale, disputed, superseded, retired, or draft	Yes
Freshness tier	Tier 0 to Tier 4 based on volatility, risk, and AI exposure	Yes
Review cadence	Monthly, quarterly, semiannual, annual, event-triggered, or continuous monitoring	Yes
Last reviewed date	Most recent human-approved review date	Yes
Next review date	Required review date before the object loses decision-grade confidence	Yes
Review triggers	Org change, system change, policy change, AI usage, incident, vendor change, legal change, owner change, metric drift, or user feedback	Yes
Evidence required	Sources needed to prove the object is accurate and still applicable	Yes
AI usage boundary	Whether AI can retrieve, summarize, recommend, classify, transform, or act using this object	Yes
Human approval rule	Who must approve changes before the object is republished or ingested by AI	Yes
Dependencies	Systems, policies, roles, data objects, decisions, prompts, workflows, or downstream artifacts that depend on this object	Yes
Change summary	What changed, why it changed, and who approved it	Yes
Confidence rating	High, medium, low, provisional, stale, disputed, or retired	Yes
Superseded by	Replacement object ID or version when this object is no longer canonical	Required when superseded

Freshness tier model

Tier	Use when	Examples	Review cadence	AI boundary
Tier 0 - Durable principle	Low volatility knowledge that rarely changes	Framework definitions, philosophy, durable principles	Annual or event-triggered	AI may summarize with citation and version metadata
Tier 1 - Operating guide	Moderate volatility guidance that changes as the organization matures	Ownership templates, decision models, meeting architecture, onboarding guides	Quarterly to semiannual	AI may retrieve and draft, but must show version and owner
Tier 2 - Decision or policy rule	Rules that affect authority, compliance, funding, workflow, control, or enterprise behavior	Approval thresholds, escalation rules, access policies, AI use rules	Monthly to quarterly plus event-triggered	AI may explain. Human approval required before action
Tier 3 - System or data object	Knowledge tied to platforms, integrations, data, metrics,	Data definitions, lineage maps, integration docs, dashboard	30 to 90 days plus system-change trigger	AI use limited to reviewed sources with freshness

Tier	Use when	Examples	Review cadence	AI boundary
	dashboards, or workflows	logic		validation
Tier 4 - AI-critical source	High-risk or high-volatility content used directly by AI agents, copilots, RAG, automation, or customer/employee-impacting workflows	Prompt libraries, retrieval sources, agent instructions, model rules, customer-facing knowledge	Continuous, biweekly, or change-triggered	AI must validate freshness and route stale sources to human review

Freshness review checklist

Review area	Question	Proof of completion
Ownership	Does the object still have a named accountable owner and knowledge steward?	Owner confirmed, steward confirmed, backup owner named
Purpose	Does the object still support a real decision, workflow, control, onboarding step, or AI use case?	Purpose confirmed or object retired
Audience	Is the audience still correct and is the content written for that audience?	Audience confirmed, irrelevant sections removed
Source of truth	Is the canonical location still correct and are duplicate copies marked as references only?	Canonical URL/system confirmed, duplicates linked back
Evidence	Are facts, claims, rules, and examples supported by current evidence?	Evidence link, owner, date, and confidence included
Freshness tier	Is the freshness tier still appropriate based on risk and AI exposure?	Tier confirmed or updated
System dependency	Have related systems, workflows, integrations, dashboards, or APIs changed?	Dependencies reviewed with system owners
Decision dependency	Have decision rights, approval paths, escalation rules, or governance policies changed?	Related decision objects checked
AI retrieval	Is this object accessible to AI? If yes, should it be?	AI usage boundary confirmed and metadata attached
Human review	Is there a clear rule for when a human must approve or reject AI use of this object?	Approval role and review trigger defined
Versioning	Are version number, change summary, approver, and supersession links updated?	Version record complete
Publication	Are website, PDF, DOCX, Markdown, JSON, and in-app copies synchronized?	All render targets updated or marked stale

Three versions by company scale

Company size	Profile	Main job	Scope	Minimum operating model	Red flags
500+ employees	Scaling company	Prevent tribal knowledge from becoming hidden operating infrastructure.	Review the top 25 to 50 knowledge objects that guide ownership, decisions, onboarding, AI pilots, customer processes, and executive reporting.	One business owner, one steward, quarterly review, simple stale/current labels, and a manual register.	Founder-memory docs, unowned playbooks, stale onboarding material, AI pilots using copied PDFs, and dashboards without definitions.
5,000+ employees	Scaled enterprise	Stop functional teams from operating with competing truth across departments.	Review the top 100 to 250 enterprise knowledge objects across functions, systems, metrics, policies, controls, and AI-enabled workflows.	Central register, functional stewards, freshness tiers, review calendar, approval workflow, and search/retrieval metadata.	Duplicate policy pages, contradictory dashboards, stale process docs, disconnected knowledge bases, and AI summaries using outdated sources.
10,000+ employees	Enterprise ecosystem	Govern knowledge as a control layer across business units, regions, systems, vendors, and AI agents.	Review all critical knowledge domains with risk-tiered governance, lineage, automated stale detection, model retrieval controls, and audit evidence.	Federated ownership, enterprise knowledge council, automated freshness signals, version lineage, AI retrieval gating, and audit-ready supersession trail.	Regional variants without governance, vendor-driven docs, uncontrolled AI retrieval, policy fragmentation, and old versions resurfacing as truth.

Review workflow

Step	What happens
1. Inventory	List the knowledge objects that matter to decisions, governance, onboarding, AI, systems, and operating model clarity.
2. Classify	Assign object type, LPM layer, owner, audience, source-of-truth location, AI exposure, and freshness tier.
3. Review	Check ownership, evidence, source validity, system dependency, decision dependency, AI usage, and publication state.
4. Score	Score the object against owner clarity, freshness, evidence, version control, AI safety, and operational integration.
5. Act	Keep current, revise, escalate, supersede, retire, or block AI retrieval until review is complete.
6. Publish	Update the website, PDF, DOCX, Markdown, JSON, in-app workflow, and AI retrieval metadata from the canonical object.
7. Monitor	Trigger future reviews from time, ownership changes, system changes, policy changes, incidents, user feedback, or AI usage.

Scoring model

Score each object from 0 to 30. 24-30 = current and AI-ready. 18-23 = usable with review notes. 12-17 = limited use. 0-11 = stale, disputed, or blocked from AI retrieval.

Dimension	What to inspect	Score
Owner clarity	Named accountable owner, steward, backup, and approval role	0-5
Freshness discipline	Clear tier, cadence, last review date, next review date, and event triggers	0-5
Evidence strength	Claims tied to current evidence, source owner, date, and confidence level	0-5
Version control	Canonical location, published formats, change history, and supersession rule are complete	0-5
AI safety	AI usage boundary, retrieval status, human review rule, and stale-source handling are defined	0-5
Operational integration	Object connects to decisions, workflows, systems, controls, onboarding, or Lapemo modules	0-5

Review register template

Object ID	Object name	Owner	Tier	Status	Last review	Next review	AI usage	Action
KFR-001	Ownership Map	Transformation leader	Tier 1	Current	YYYY-MM-DD	YYYY-MM-DD	Retrieve + summarize	Keep
KFR-002	AI Initiative Owner Register	AI governance owner	Tier 2	Under review	YYYY-MM-DD	YYYY-MM-DD	Retrieve only	Update owner
KFR-003	Customer policy FAQ	Customer ops owner	Tier 4	Stale	YYYY-MM-DD	YYYY-MM-DD	Blocked	Escalate
KFR-004	Revenue dashboard definition	Data product owner	Tier 3	Disputed	YYYY-MM-DD	YYYY-MM-DD	No AI use	Resolve metric
KFR-005	Meeting Architecture	Operating model owner	Tier 1	Current	YYYY-MM-DD	YYYY-MM-DD	Draft support	Publish

AI prompts for future Lapemo skill

Prompt	Instruction
Freshness classifier	Classify this object by freshness tier based on volatility, risk, decision impact, and AI exposure. Explain the tier and review cadence.
Staleness detector	Review this object metadata and identify missing owner, expired review date, weak evidence, stale source, unresolved dependency, or unsafe AI usage.
Review assistant	Generate a review checklist for this object based on its LPM layer, audience,

Prompt	Instruction
	source-of-truth location, and downstream uses.
Supersession assistant	Compare the prior and current version. Identify what changed, who must approve it, and which downstream artifacts need to be republished.
AI retrieval guard	Determine whether this object can be retrieved by AI. If not, explain the metadata, evidence, review, or approval gap that blocks usage.

Validation rules

Rule	Standard
No owner, no publication	A knowledge object cannot be published or used by AI without an accountable owner.
No review date, no confidence	Any object without last reviewed and next review dates is provisional at best.
No evidence, no decision-grade claim	Claims used for decisions, governance, or AI recommendations require evidence metadata.
No stale AI source	Objects marked stale, disputed, superseded, or retired must be blocked from AI retrieval unless explicitly approved for historical reference.
No duplicate truth	Published copies must reference the canonical source-of-truth object and version.
No silent supersession	New versions must state what changed and which older objects are replaced.

How this becomes reusable

- Human artifact: downloadable DOCX and PDF for workshops, consulting, and website resources.
- Website artifact: Markdown page that explains the method and lets teams copy the template.
- Machine-readable object: JSON schema containing fields, scoring, prompts, validation rules, render targets, and ingestion logic.
- Lapemo skill: guided workflow that inventories knowledge objects, classifies freshness tier, flags stale content, routes review to owners, and updates publication outputs after approval.
- AI boundary: AI can help detect staleness, compare versions, draft updates, and route approvals, but cannot silently mark high-risk knowledge current without human approval.

Recommended publication block

Use the Knowledge Freshness Review to keep operating-model artifacts, AI-readable content, policies, dashboards, prompts, and decision guides current. This template helps teams assign ownership, set review cadence, validate evidence, control AI retrieval, and retire stale knowledge before it becomes enterprise truth.