AutoRFP.ai

AUTORFP.AI RESEARCH

INTERACTIVE WORKBOOK

Content Library Maturity Checklist

Is your content library ready for AI?

Seven levels separate the teams that win on automation from the teams stuck searching. Find your level, tick what is true today, and read your gaps as the work ahead.

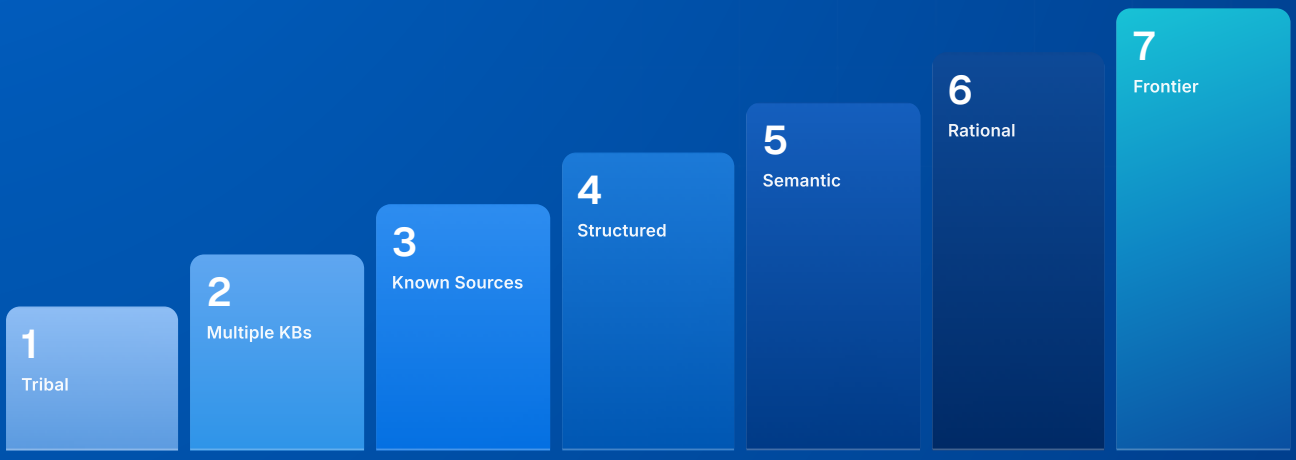

FORMAT

Self-assessment · tick as you go

TIME

15 minutes

autorfp.ai

HOW TO USE THIS

There are seven levels of content library maturity

1. **Find your level.** Read the seven stages below and the detail pages that follow.
2. **Tick what is true today.** Boxes save as you go — return any time.
3. **Read the gaps.** Unticked boxes are the work in front of you.

OVERVIEW — THE SEVEN LEVELS AT A GLANCE

LVL	STAGE	WHAT IT LOOKS LIKE
1	**Tribal**	In people's heads and docs. No formal places.
2	**Multiple KBs**	Notion, Drive, SharePoint. Duplicated and drifting — five docs, three disagree.
3	**Known Sources**	Canonical home, named owner, review cadence. One place for each topic.
4	**Structured KBs**	Categorised and tagged across multiple dimensions. Contextual and easy to find.
5	**Semantic KBs**	Embedded for AI. Search by meaning, not keywords. Answers in seconds, truth-checked.
6	**Rational KBs**	A reasoning layer weighs recency, authority and context — and answers before you ask.
7	**The Frontier**	What comes after rational — a library that learns from every answer. Where this goes next.

THE GAP, IN ONE LINE

59% of top-performing bid teams use content library automation. Only **36% of low-performing teams** do.

1 LEVEL 1 Tribal

0 / 3

Knowledge lives in people, not places. Tick the boxes that are true for you — if most are, you are at Level 1, and the next move is getting content out of individuals' heads into shared systems.

- [ ] Answers live in individuals' files, inboxes, and old proposals. There is no agreed place to look first.
- [ ] New joiners ask a person, not a system. Knowledge walks out the door when people do.
- [ ] The same question gets answered differently depending on who happens to respond.

CLIMB Write it down and put it somewhere shared.

2 LEVEL 2 Multiple KBs

0 / 3

Content is written down but scattered across tools, duplicated, and drifting. Five docs, three disagree. Tick every box to confirm Level 2.

- [ ] Content is written down and held in **shared systems** (e.g. Notion, Drive, SharePoint), not in personal files.
- [ ] Text is **machine-readable and extractable**, not locked inside scanned PDFs, images, or proprietary formats.
- [ ] A new joiner can find an answer without DMing a colleague, even if it takes a few clicks across tools.

CLIMB Agree one canonical home per topic and give it an owner.

LEVEL 3

3 Known Sources

0 / 12

Each topic has a single home, a named owner, and a review cadence. People know where to look and trust what they find. Tick every box to confirm Level 3.

- [ ] **One canonical home per topic.** Everyone knows the single place to look for a given subject.
- [ ] **Named owner per item.** A specific team or individual is responsible — documented, not assumed.
- [ ] **Ownership is visible in the system.** Anyone can look up who owns an item without asking around.
- [ ] **Ownership sits with a team**, not only an individual, so it survives someone changing roles or leaving.
- [ ] **A reassignment process exists.** When someone leaves, their content gets reassigned. No orphaned answers.
- [ ] **A source of truth is defined** for cross-system content. One system is master and everyone knows which.
- [ ] **Content age is tracked and visible.** You can see when each item was last reviewed, not just created.
- [ ] **A review cadence exists and is enforced.** Quarterly, biannually, annually — having one and holding it.
- [ ] **Cadence is set per topic**, not one blanket rule. Fast-moving content is reviewed more often.
- [ ] **Expiring content has a fixed date.** Time-bound knowledge is refreshed or retired on schedule.
- [ ] **Stale content is flagged or archived.** Items past their review date get surfaced, not left to rot.
- [ ] **Review & approval workflows exist.** Regulated content routes through compliance; technical through engineering.

CLIMB Structure it. Add categories, types, tags, and dimensions.

LEVEL 4

4 Structured KBs

0 / 12

Content is categorised and tagged across multiple dimensions, so the right answer is fast to find and easy to maintain. Tick every box to confirm Level 4.

- [ ] **Organised by folder or category** with a logical structure a new joiner could navigate unguided.
- [ ] **Split into defined types:** capability statements, compliance, case studies, specs, pricing — each labelled.
- [ ] **Tagged with metadata:** product, region, persona, use case, industry. Tags make content findable.
- [ ] **Tags & folders follow a documented taxonomy** — written down somewhere, not in one person's head.
- [ ] **Modelled across multiple dimensions**, treating Region and Product as separate axes, not one rigid tree.
- [ ] **A single item can sit in many folders/tags (1:many)**, tagged across all of them rather than copied.
- [ ] **Hierarchies go as deep as needed** (Global → North America → US → California), so nuance is not flattened.
- [ ] **Sensitive content has restricted edit access.** Regulated, legal, compliance — only authorised people edit.
- [ ] **Cross-system content syncs automatically.** Updates flow downstream — no one copy-pastes between tools.
- [ ] **Content usage is tracked.** You know which items get used, how often, and in which bids.
- [ ] **Unused content is identifiable.** You can pull items untouched for 6+ months and decide what to do.
- [ ] **Most-used content is identifiable** and gets priority review and the most attention during updates.

CLIMB Make it semantic and machine-accessible.

LEVEL 5 · AI-READY

5 Semantic KBs

0 / 8

Content is embedded for AI. Retrieval is by meaning — auditable, ranked, measured, and checked against the source, so the system reasons about what is true rather than passively storing it. Tick every box to confirm Level 5.

- [ ] **Embedded in a vector or graph database.** AI finds relevant content by meaning, not just keyword matching.
- [ ] **Reachable via MCP or API**, so AI agents and external tools can use the library programmatically, not only through a UI.
- [ ] **Every item carries instructions for AI use.** Per-item or system rules cover tone, formatting, and what to prioritise.
- [ ] **AI search and retrieval is auditable.** You can see what was pulled, why, and trace each answer back to its source.
- [ ] **Content ranking is transparent.** You understand why a given item surfaces first.
- [ ] **Prioritisation is configurable** by item, type, and folder/tag — so quality, recency, and relevance shape what AI surfaces.
- [ ] **Retrieval is measured** with an evaluation framework. AI grading or human review tracks whether the right content surfaces.
- [ ] **The system reasons about what is true**, not just what is stored. Answers are checked against the source, not confidently stale.

TWO STAGES STILL SIT ABOVE

Level 5 makes AI **cite instead of guess**. Above it sit the **Rational KB** that reasons about what's true, and the frontier beyond that. Turn the page.

LEVEL 6 · THE 2026 LEVEL

6 Rational KBs

0 / 6

A reasoning layer sits over your metadata and works out what is true — automatically weighing recency, authority, and context for whoever is asking, and composing the answer before the question is even asked. Tick every box to confirm Level 6.

- [ ] **A reasoning layer sits over your content**, deciding what is true rather than returning the closest raw match.
- [ ] **Conflicts resolve by recency.** The system knows which source is the most recent.
- [ ] **Conflicts resolve by authority.** It knows which source is the most authoritative.
- [ ] **Conflicts resolve by context.** It knows which source is most relevant to the question being asked.
- [ ] **Answers adapt to who is asking**, in what context, and when.
- [ ] **Answers are pre-composed**, combining time and context before the question is even asked.

LEVEL 7 · THE FRONTIER

7 What comes after rational?

Beyond a library that reasons is one that **improves itself** — learning from every answer, closing its own gaps, and anticipating the next question before anyone asks it. No one is fully here yet. It is the frontier of the curve, and where we think this goes next.

BONUS WORKSHEET

Designing your hierarchy

Levels 4 and 5 depend on a hierarchy that does not collapse under its own weight. Most libraries fail here — one rigid folder tree, the same answer copied into ten places, detail flattened that mattered. Four design principles fix that.

1

Think in dimensions, not one hierarchy

Product, Region, Persona, and Content Type are separate axes. Combine them at retrieval instead of multiplying them into one folder path. Twelve folders like "Contracts UK" and "Insurance US" collapse into two clean dimensions — roughly half the options, clicks, and training.

2

Let one item live in many places (1:many)

A single answer often belongs to several folders or tags at once. Tag it across all of them rather than copying it — so there is one item to maintain, not five copies to drift apart over time.

3

Give the hierarchy depth where it earns it

Some branches need more levels for nuance (Global → North America → United States → California). Add depth only where the content demands it, and keep the rest shallow.

4

Set ownership, cadence & expiry per branch

Assign each area to a team rather than an individual, so it survives people moving on. Pick a review cadence that fits the topic (annually for stable, monthly for fast-moving), and put a hard expiry date on knowledge that goes out of date.

Work through the four principles on the page opposite — or hand them to an AI assistant using the prompt overleaf, with your own context filled in.

THE PROMPT

Turn the four principles into a structure

Fill in the bracketed fields and paste this into your AI assistant. It turns the four design principles into a structure you can build.

HIERARCHY DESIGN PROMPT

You are helping me design the content hierarchy for my RFP and proposal content library. I want a structure that is easy to navigate, scales without folder sprawl, and is ready for AI retrieval.

My context:
- What we sell: **[products or platforms]**
- Where we sell: **[regions or markets]**
- Who we sell to: **[personas or industries]**
- Content types we hold: **[capability statements, compliance, case studies, security, pricing, technical specs]**
- Regulated or restricted content: **[anything that needs controlled access or compliance sign-off]**

Design my hierarchy using these four rules:
1. Dimensions, not one tree. Treat Product, Region, Persona, and Content Type as separate dimensions I can combine, instead of one path like Product > Region > Type that multiplies into dozens of near-duplicate folders.
2. One item, many places. Identify content that belongs in more than one location (e.g. a case study relevant to two products and three regions) and tag it across all of them rather than duplicating it.
3. Depth where it earns it. Add extra levels only where the content needs the nuance (e.g. Global > North America > United States > California). Keep shallow branches shallow.
4. Ownership, cadence, and expiry per branch. For each top-level area tell me which team should own it, how often it should be reviewed, and whether any of it expires on a fixed date.

Output: a proposed structure I can build, a list of items that should be tagged across multiple dimensions, and any places where my current setup would create duplication.

PREFER TO WORK IT THROUGH YOURSELF?

Four moves to a structure that scales

0 / 4

No AI required. Run the four principles by hand, in order. Use the space under each to sketch as you go.

- [ ] **List your dimensions.** Aim for three to five axes, not one mega-tree.
- [ ] **Mark the content that needs to live in more than one place.** That is your tagging list — not your duplication list.
- [ ] **Mark where you need more depth, and where you are over-engineering.** Add levels only where nuance is real.
- [ ] **Assign each branch a team owner, a review cadence, and an expiry rule.** Team over individual. Cadence to match the topic. A date for anything that expires.

WHERE THIS LANDS

Your unticked boxes across all seven levels are your roadmap. Start at the lowest level with a gap and climb one rung at a time — structure before semantics, always.

AutoRFP.ai

BUILT WITH EVERY FEATURE ON THIS LIST

Climb the curve faster than your competition.

AutoRFP.ai is the Rational Knowledge Base in one platform — named owners, per-topic cadences, multi-dimensional hierarchies, semantic retrieval, and a reasoning layer that weighs recency, authority, and context. And we keep building the next level, so you stay ahead.

See it on your own content — book a demo

WEB	RATED ON G2	DOCUMENT
autorfp.ai/book-demo	4.9 / 5 · 982+ reviews	AutoRFP.ai Research · 2026