

The Topical Map and Content Cluster Workbook: Seeds, Page Roles, Links and Refresh

A planning workbook for a topical map: a seed-to-cluster method, a page-role sheet, an internal linking plan and a refresh cadence, with a plain statement of what is and is not known about topical authority.

Contents

01 What a topical map is, and what the evidence says about topical authority

02 From seed to cluster: the method, step by step

03 Page roles and the page-role sheet

04 The internal linking plan

05 Refresh cadence: review by how fast the claims go stale

06 Putting the workbook to work: a first month

Key Takeaways & Sources

01

What a topical map is, and what the evidence says about topical authority

A topical map is a planning document that lays out the subjects a site intends to cover, the questions under each subject, the single page that answers each question, the role each page plays and the links that connect them. It exists to prevent two common failures. One is gaps, where a buyer asks a question at a key stage and the site has nothing. The other is overlap, where five pages half-answer the same question and none is the one worth ranking. A map makes both visible before anyone writes.

The industry word for the hoped-for result is topical authority, the idea that a site which covers a subject thoroughly and consistently is treated as more credible on it. We want to be careful with that term. As of 7 October 2026, none of the Google Search Central pages we read while preparing this workbook (the SEO Starter Guide, the people-first content guidance, the link and canonicalization guides, the AI features guide, Search Essentials and the spam policies) contains the phrase. That is a statement about those pages and not about everything Google has said or about how its systems work. It means we cannot cite Google for topical authority as a defined, measurable thing.

What Google does document is a set of questions about content. Among them: does the content provide original information, reporting, research or analysis; does it offer a substantial and complete description of the topic; and does the site have a primary purpose or focus. Those are quality questions, and a topical map helps you answer them honestly, because it shows what the site covers, in how much depth and with what overlap. Use the map for that purpose. Do not report a map as proof that authority has been gained.

The same discipline applies to AI answer surfaces. Google states that there are no additional requirements to appear in AI Overviews or AI Mode, no special optimizations are necessary and no special schema.org markup is needed; a page must be indexed and eligible to appear in Search with a snippet. So planning a cluster for AI answers is the same work as planning it for search: one clear page per question, written to be useful, crawlable and linked. Nothing in a topical map is a separate AI tactic.

Finally, a map is not a mandate to publish a lot. The people-first guidance warns against creating content primarily to attract search visits. A cluster of eight pages that each answer something real beats forty pages produced to fill a grid, so dropping a candidate page that adds nothing is part of the work.

That is a statement about those pages and not about everything Google has said or about how its systems work.

02

From seed to cluster: the method, step by step

The method has seven steps and each produces something you can inspect. Begin with seeds. A seed is a subject the business could credibly teach, taken from what it sells and what its customers need help with, and written as a short noun phrase. Three to six seeds is plenty for a first map. A seed list taken straight from a keyword tool tends to describe what other sites cover, not what your buyers ask your sales team.

Next, harvest questions under each seed. Sources, in order of how much we trust them: questions customers asked on sales calls and in support, the search terms people use on your own site, queries from your search console report, and the questions that appear in search results for the seed. Record the wording exactly as asked and the source of each. A question without a source is a guess, and the map should show which parts are guesses. Keyword tools can add volume estimates, and those are estimates from the tool vendor, not counts of searches. Note the tool and the date next to any figure you keep.

Then group the questions by the reader's job, not by vocabulary. "What is X", "how does X work" and "why does X matter" may be one page for a reader who is learning, while "X versus Y" and "X pricing" belong to a reader who is deciding, and "how to set up X" belongs to a reader who has already chosen. Group by the stage and the task, and name each group with the job in a few words. This is where most maps change: groups that looked separate in a keyword list turn out to be one reader asking one thing three ways.

Check each group for overlap with what already exists. List the current pages that touch the group, mark the one best positioned to serve it and decide for the others. Options are to merge into the best page and redirect, to rewrite so the page serves a different intent, or to leave it if the overlap is only partial. This step prevents a new page being created for a question the site already answers. Google's canonicalization guidance describes consolidating signals for duplicate or similar pages into one preferred URL as one reason to specify a canonical; use that when two near-identical URLs must both exist.

Now assign roles and decide what is not worth a page. Every surviving group gets a role from the table in the next section, a working title and a primary question. If a group cannot name a role or a question, drop it and record why. Finish by drawing the links, which are covered in the linking section, and by setting review dates, covered in the cadence section. The output is the page-role sheet, which is the map in a form you can sort and filter.

Step	Input	Output	Gate to pass before moving on
1 Seeds	The offers, the customers' problems	Three to six seed subjects as short noun phrases	Each seed is something the business could credibly teach

Step	Input	Output	Gate to pass before moving on
2 Questions	Sales and support notes, site search, search console queries, search results	A list of questions, exact wording, each with a source	Every question has a source or is marked as a guess
3 Groups	The question list	Groups named by the reader's job and stage	Each group is one task; questions that mean the same thing are together
4 Overlap	The group list and a crawl or sitemap of current pages	For each group: best existing page, or none, and a decision for any other page	No group is served by two pages without a recorded decision
5 Roles	The groups that survived	Role, working title and primary question for each page	A group that cannot name a role or question is dropped with the reason noted
6 Links	The page-role sheet	A linking plan: who links to whom, with anchor text	No planned page is without a route in from another page
7 Refresh	The page-role sheet	A review date and a volatility class for each page	Every page has a next review date and an owner

03

Page roles and the page-role sheet

A role says what a page is for in the cluster. Naming roles stops every page trying to be everything, and it makes the linking plan almost write itself, because a role implies where the page should send the reader next. The table lists the eight roles we use. They are our working categories, not a standard, so rename or merge them to suit the site.

The page-role sheet has one row per page, whether planned or existing. Its columns are the cluster, a page identifier, a working title, the primary question in the reader's words, the reader's job, the role, the intent (learn, compare, decide or do), the target URL, the owner, the status and an overlap check. The overlap check is a short note stating which other pages were looked at and why this one is distinct. Writing that note is what gives the check force. A blank cell means the check did not happen.

Apply one rule hard: one primary intent per page. If the page title needs "and" to describe what it answers, it is probably two pages. If two rows in the sheet have the same primary question, one of them is wrong. The fix is to merge the pages, to change one so that it answers a different question or, where both URLs must stay live, to point one at the other with a canonical. Do not leave the overlap in place on

the hope that search engines will sort it out. They may, but you will not be able to tell which page they chose or why.

Track status with a small set of values: planned, drafted, published, indexed, review due and retired. Two ratios are worth computing. Coverage is the number of published pages divided by the number of planned pages in the cluster. Overlap is the number of questions served by more than one page, divided by the number of questions. Both are simple counts of your own sheet. They show progress against your plan, which is all they can show, and neither is a measure of search performance.

Role	Reader's job	What the page does	Links to
Hub	Get oriented in the subject	Frames the topic, lists the questions and points to each page that answers one	Every supporting page in the cluster
Definition	Understand a term	Defines one term plainly and says where it is contested	The hub and the related how-to
Explainer	Understand how something works	Explains a mechanism with evidence and limits	The hub, the related decision page
How-to	Do a task	Steps with checks and what can go wrong	The tool or template, and the next task
Comparison	Choose between options	States the criteria, the options and what would change the answer	The hub and the commercial page, where it fits
Proof	Check that it works	A case, an example or a method with data and limits stated	The relevant explainer and the commercial page
Tool or template	Get a working file	A downloadable worksheet with instructions	The how-to that explains it
Commercial	Take the next step with the company	Describes the offer and how to start	Proof and the hub

04

The internal linking plan

Internal links do two jobs. They help people find the next useful page, and they help search engines discover pages and understand what each one is about. Google's link guidance says Google can generally crawl a link only if it is an anchor element with an href attribute, and that most links in other formats are not parsed by its crawlers. So the first check in the plan is mechanical: every link that matters is a real anchor with an address. A link built with a script that sets the destination on click may work for visitors and be invisible to a crawler.

The second check is anchor text. Google says good anchor text is descriptive, reasonably concise and relevant to both the page it sits on and the page it links to, and that for an image used as a link the alt text serves as the anchor text. "Click here" tells the reader and the crawler nothing. The plan has a column for the exact anchor text, so a reviewer can see what each link says before it goes live. Clear, accurate text is enough.

The pattern we start from is simple and based on the roles. Each supporting page links up to its hub and onward to the next step in the reader's task. The hub links down to every supporting page in the cluster. Pages whose roles are adjacent, such as a how-to and the template it uses, link to each other. Comparison and proof pages link to the commercial page where it fits the reader's next step. Links sit in the body at the point where the reader would want them, not in a block at the foot of the page. We do not set a minimum or maximum count, because the right number depends on the page, and any number from us would be invented.

Test for orphans. An orphan is a published page with no internal links pointing to it, which leaves a crawler to find it only through a sitemap and leaves readers without a route. Crawl the site, export the inbound link counts and list the pages with none. Each orphan gets a decision: link it from the right place, merge it, or retire it. Run the same check on pages planned but not yet written, by asking which existing page will link to the new one on the day it publishes.

The SEO Starter Guide adds a note of proportion that is worth keeping in view. It suggests organizing a site in a logical way, so that search engines and users can understand how pages relate, and also says not to drop everything to reorganize, because search engines will likely understand pages as they are now. A linking plan is an improvement program, not a rebuild. Fix the orphans and the obvious dead ends first, and add the cluster links as pages are written or revised.

- Every link that matters is an anchor element with an href attribute, checked in the rendered page
- Anchor text is descriptive, concise and relevant to both pages; images used as links carry useful alt text
- Hub links to every supporting page; each supporting page links to the hub and to the next step in the reader's task
- No orphans: every published page has at least one internal link in from a relevant page, or a recorded decision

The plan has a column for the exact anchor text, so a reviewer can see what each link says before it goes live.

05

Refresh cadence: review by how fast the claims go stale

A map decays. Products change, laws change, platforms rename things and a statistic that was current when a page was written is quietly out of date a year later. The cure is a review schedule that depends on how fast each page's claims go stale, not on one blanket rule such as "update everything yearly". Classify each page by the most volatile claim it makes, and review it on that class's schedule.

The table gives four classes with starting intervals. Those intervals are our suggestions as a starting point, not findings, and you should shorten them if a page has been wrong before or lengthen them if reviews keep returning "no change". Add a trigger for each class too. A review should also run when an event makes the page suspect, such as a change to the product, a regulator's notice or a platform's documentation update, whether or not the date has come.

At each review, decide among five outcomes and record which one. Keep means the page is accurate and still the right page. Update means a claim, an example or a source needs changing. Merge means the page now overlaps another and should be combined with it. Redirect means the page is replaced by a better URL. Remove means the page has no remaining purpose and should be taken down with the right status code. Record the reviewer, the date and the outcome in the cadence sheet so the history is visible.

Be strict about dates on the page. If you show a reviewed date, change it only when someone has actually checked the claims. Moving a date without checking is a small dishonesty that readers notice and that the next reviewer cannot detect. Where you cannot confirm a claim any longer, take it out and say so in the change note rather than leaving it standing on the strength of an old source.

Volatility class	Examples of claims	Starting review interval (our suggestion)	Review also when
Durable	Definitions, principles, how a method works	About every twelve months	A source it relies on is retired or revised
Tool-dependent	Steps in a product interface, settings, features	About every six months	The tool announces a change or a step stops working
Rule-dependent	Laws, platform policies, guidance from a regulator	About every three months	A notice or an amendment is published
Data-dependent	Statistics, prices, benchmarks, dated examples	About every three to six months	A newer edition of the source appears

06

Putting the workbook to work: a first month

Start small. Pick one cluster, the one closest to revenue, and run the seven steps for it alone. A first cluster of eight to twelve pages is enough to learn how the sheet behaves. It also shows how long the overlap check takes, which is usually longer than expected.

In the first week, collect seeds and questions. In the second, group them and check overlap against a crawl of the current site. In the third, assign roles, draft the page-role sheet and the linking plan, and have someone outside the team read it. In the fourth, set review dates, brief the first pages and agree who owns the sheet.

Decide how you will know the work is doing something, and decide it before you start. Page counts and coverage tell you about progress against the plan. Evidence about outcomes has to come from elsewhere: whether the pages are indexed, whether queries from the question list start to appear in your search console report, whether sales hears fewer repeated questions and whether the pages are used in the sales process. None of those is proof of authority, and some take months to show. Note the baseline for each at the start so a change can be recognized later.

Last, protect the map from its own success. Once it exists, there will be pressure to add pages because a keyword has appeared. Apply the same gate as before: a source for the question, a role, a distinct intent and an owner. If a request cannot pass, record it and move on.

Week	Work	Gate to pass before moving on
1	Choose the cluster; write the seeds; harvest questions with sources	Every question has a source or a guess flag
2	Group questions by reader job; crawl and check overlap	Every group has a decision on existing pages
3	Assign roles; draft the page-role sheet and the linking plan; outside read	A reader outside the team has listed questions with no page
4	Set review dates and owners; brief the first pages; record baselines	Every page has an owner and a next review date; baselines written down

Key Takeaways

- A topical map is a plan that assigns each reader question to one page, gives every page a role, and says how the pages link. It is a planning document, and it is only as good as the questions it starts from.

- Topical authority is an industry term. None of the Google Search Central pages we reviewed for this workbook defines or uses it, so treat it as a working aim for coverage and consistency, not as a score you can read or a ranking factor Google has described.
- Start from the business offer and the questions buyers ask, not from a keyword export. Group by the job the reader is trying to do, then check each group for overlap before any writing is commissioned.
- One primary intent per page. Two pages that answer the same question compete with each other, and the cleanup options are to merge them, to separate their intents or to consolidate the duplicate with a canonical.
- Google documents that links are crawled when they are anchor elements with an href attribute, and that descriptive anchor text helps people and Google understand the target. The linking plan is built on those two facts.
- Put a review date on every page, set by how quickly its claims go stale, and decide at each review among keep, update, merge, redirect or remove. Record the decision.

Sources

1. SEO Starter Guide: The Basics — Google Search Central —
<https://developers.google.com/search/docs/fundamentals/seo-starter-guide>
2. Creating Helpful, Reliable, People-First Content — Google Search Central —
<https://developers.google.com/search/docs/fundamentals/creating-helpful-content>
3. SEO Link Best Practices for Google — Google Search Central —
<https://developers.google.com/search/docs/crawling-indexing/links-crawlable>
4. How to Specify a Canonical with rel="canonical" and Other Methods — Google Search Central —
<https://developers.google.com/search/docs/crawling-indexing/consolidate-duplicate-urls>
5. AI Features and Your Website — Google Search Central —
<https://developers.google.com/search/docs/appearance/ai-features>