BuildUp BOOKINGS

01

BONUS 01 OF 11 · MASTERING VACATION RENTAL MARKETING

AI Search & GEO Readiness Checklist

A companion resource to the 2nd edition

WHAT YOU'LL HAVE WHEN YOU FINISH

A scored, dated baseline of how visible your company is to AI search tools, and a ranked list of the specific fixes that will move it.

INSIDE

01 Do this now
02 The depth
03 Worked example
04 Your turn
05 How you'll know it worked
06 Where this lives in the book

By Conrad O'Connell
BuildUp Bookings

Facts checked November 2026
buildupbookings.com/book-bonuses

PART 1

Do this now

Start here. Each one works on its own.

Eight actions. Nothing here requires a developer, a budget, or a vendor. If you read no further than this page, do these.

1 **Count the H1 tags on your homepage.**

Open your homepage, right-click, View Page Source, search for `<h1`. There should be exactly one. If your booking widget is rendering CHECK-IN, CHECK-OUT and BEDROOMS as headings, every crawler that visits believes those are your page's main topics. This is usually a one-line template fix and it is the cheapest meaningful work in this document.

2 **Confirm you are not blocking the AI search crawlers.**

Go to `yourdomain.com/robots.txt` and read it. Then check your CDN and your host separately — Cloudflare and others ship AI-bot blocking that can be switched on by default, without anyone deciding to turn it on, and some managed WordPress hosts block AI crawlers at the server, whatever robots.txt says. Being invisible because a security default fired is the worst way to lose this. While you are in Google Search Console, open Settings, then Search generative AI, and confirm it still says Include.

3 **Test whether your pages actually emit schema.**

Paste your homepage and one property page into Google's Rich Results Test, then into the Schema Markup Validator at validator.schema.org, which shows every type you emit, including ones Google no longer displays as rich results, such as FAQPage. Having an SEO plugin installed is not evidence that it is working — page builders strip schema output regularly, and nobody notices for years.

4 **Create a Wikidata entry for your company.**

Free, roughly an hour, and it gives your company a public, structured record of its basic facts that other databases and tools reuse. Founded date, location, business type, markets served. Wikidata keeps only items backed by serious, publicly available references, so have two or three ready (press coverage, your state business registration) and disclose that you work for the company. No AI company documents using Wikidata to recognize businesses, so treat it as a cheap bet, not a proven lever.

5 **Put a real human byline on your content.**

Not “The Team.” A named person, a photo, and two lines on why they know this market — how long they have lived there, which beach is better in October.

6 **Write ten prompts and run them.**

Long, conversational questions a real guest or homeowner would actually type into ChatGPT. Run each one several times across two or three tools. Write down how often your company gets named. Date it. This is your baseline, and without it you cannot tell whether any of this worked.

7 **Add one answer block to your single best page.**

A short, complete answer, a paragraph or two, sitting directly under a heading that asks the question, answering it completely, making full sense pasted into a blank document with nothing around it. Answer first. Sell second. You will see “134 to 167 words” quoted as the ideal length; I have not found a study behind it, and Google says there is no ideal length.

8 **Look yourself up where AI looks for lodging.**

Find your company on your convention and visitors bureau’s lodging page, your state tourism site’s places-to-stay page, and the sites that rank local managers (AirROI, Comparent, Airbtics, AirDNA, Manager Review). Then open your Google Business Profile. Get listed where you are missing, under the name guests know you by, and correct what is wrong where you are listed. These are the pages AI reads when a guest asks where to stay.

PART 2

The depth

The full method.

The two numbers that matter

Before any tactics, get the scoreboard right. There are two things worth measuring here, and if you are tracking something else you are tracking the wrong thing.

Citations. Does an AI tool link to your site when it answers a question in your market?

Brand mentions. Does an AI tool say your company's name when someone asks who they should book with?

These are different, and the second is probably worth more. A citation sends you traffic. A mention sends you a guest who then searches your company name and books — which your analytics will record as direct traffic or branded search, and which you will almost certainly credit to the wrong channel. Hold onto that, because it is the whole measurement problem in one sentence.

And one does not reliably produce the other. In one market we audited in November 2026, the most-cited rental manager in town was named in one of 42 ChatGPT answers to questions real guests ask. A competitor whose own website had six citations was named in all 18 answers about family trips, ski weekends and the town's winter festival, because the local visitor bureau and the state tourism site both list it. The guide earns the citation. The listing earns the recommendation.

Start with a baseline, not with fixes

Almost every article written about AI search skips this step, which is how operators end up doing six months of work with no way to tell whether any of it did anything.

Prompt research is the new keyword research. Same instinct, different unit. In ordinary SEO you research what guests type. Here you research how they ask. The difference is that no AI platform publishes prompt volume. A few tools now sell modeled estimates built from consumer panels, but they are priced for enterprise and thin for a single vacation market. So you build the list yourself, out of what you already know about who books with you.

Write out the questions a real person would ask. Long ones. Conversational ones. They look nothing like keywords:

> "I'm planning a family trip with my husband and two kids. We live in Atlanta. I want a mountain escape within about three hours' drive, in a cabin that works for young children. Who should I look at and who should I book through?"

> “Who is the best property manager to work with in [your market]? I want to maximize the revenue I get from my vacation home there.”

Build ten to twenty. Cover your guest types, and if you manage for owners, build a set for them too.

Here is the part worth doing even if you never test a single prompt: writing this list exposes whether you actually understand your guest. If you cannot produce ten specific questions a real person would plausibly use to find a company like yours, that is not an AI search problem. That is a positioning problem, and it has been quietly costing you in every channel you run.

Then measure properly. Run each prompt repeatedly, across several tools. Not once. The answers vary between runs, so a single test tells you close to nothing, and building a strategy on one result is how people reach confident wrong conclusions. In our own October 2026 test, when a manager showed up for a ChatGPT prompt, it showed up in all five runs only half the time. What you want out of it is a rate: out of the times you asked, how often did your company get named? Write it down with the date. Re-run it monthly while you are actively working on this, quarterly at minimum.

Do not rely on an SEO platform’s built-in AI question set for this. Those sets track general questions, and in our November 2026 audits they held almost no booking questions for a single vacation market. In one, the platform showed zero ChatGPT citations for a site our own prompt list found cited in 9 of 48 answers. Your list is the measure.

Then look at your analytics. Since May 2026, GA4 has sorted visits from ChatGPT and similar assistants into an “AI Assistant” channel of its own. Treat it as a floor: visits that arrive without a referrer land in Direct, and the guest who read your name in ChatGPT and then Googled you shows up as branded search. Watch branded search impressions in Search Console next to it.

Set your expectations on timing. Two clocks run here. ChatGPT search, Perplexity and Google’s AI features answer from live retrieval, so a fixed page can start showing up once their crawlers fetch it again. What a model “knows” about you without searching changes only when it is retrained, on a schedule nobody publishes, and your market’s content is a rounding error in a corpus that size. Make the improvements, keep measuring, and do not conclude at the thirty-day mark that it does not work.

Winning Google does not win ChatGPT

If you have done the SEO work in the book, your instinct will be that AI search sorts itself out — that ranking well means getting cited. That is half true, and the half that is false will cost you.

For **Google’s own AI surfaces**, ranking is the foundation. A page has to be indexed and eligible to show in Search before an AI Overview can use it. But Google says AI Overviews and AI Mode may fan a question out into several related searches, so the pages they cite are not always the page-one results for the exact query the guest typed. Every bit of SEO work you have already done compounds into that visibility.

For **ChatGPT and Perplexity**, ranking is not enough. They are not simply re-reading Google's page one. They are checking whether the world already talks about you — sites that rank local managers, visitor bureaus, review platforms, press. In our October 2026 test across 84 US markets, ChatGPT read or linked a site ranking local managers in 418 of 420 answers about the best rental companies in a market, while forum and social links showed up in 2 of 1,260 answers. Ask the same question in ChatGPT and in Google and you will often see different sites cited.

That is a different problem than SEO solves. SEO asks: is my page the best answer to this query? Conversational AI visibility asks: does this company exist as a known entity in the places I trust? Ranking helps: in that test, managers in Google's top 20 for their market were named in a median 64% of AI answers, against 41% for those outside it. But you can rank first in your market and still be invisible to ChatGPT, because ChatGPT is drawing on different sources.

The Five Levers

Score yourself honestly on each of these before you touch anything. The weights are my judgment of how much each one matters, not a measured effect, because nobody has that data yet.

Lever 1 — Citability (25 points)

Can a machine lift one self-contained chunk of your page and use it as an answer, without needing the rest of the page for context?

The heaviest lever, and the one almost nobody gets right. These engines do not cite pages. They cite passages. If your answer to "when is the best time to visit" is scattered across four paragraphs interleaved with booking calls-to-action, there is no passage to lift.

Four things reliably kill citability:

Fragmentation. Page builders love to scatter text across dozens of small boxes. A page can look beautiful and contain nothing extractable. Watch for text concatenated together without separation — a phone number run directly into prose is a phone number no machine will ever parse.

Unlabeled numbers. Big design-element figures with no units and no labels. "150+" and "728K" floating in a hero section tell a reader nothing and tell a machine less. An unlabeled number is not a claim, it is decoration. Label it or lose it.

No sources and no named author. If your content makes a claim about seasonal demand and cites nothing, and no human being is named as the author, you have handed the engine every reason to prefer someone else.

Stale dates. Your event guides are probably among your most-cited pages. In both client audits we ran in November 2026, a guide to the area's big seasonal event was being cited by Google's AI Overviews while its title still carried last season's date. Update the same page, at the same address, before the season starts every year. A new page starts over; the old one already has the links and the citations.

There is a paradox worth sitting with: to do well with machines, you have to make your website more human. A named local who demonstrably lives in the market beats an unattributed post from a content team. I will be straight that this is a working theory rather than a proven mechanism — plenty of unattributed content ranks fine. But in a market with a hundred articles about the same attractions, I think the one written by a named local wins the citation. That is a bet, and I would rather tell you it is a bet than dress it up as a certainty.

Lever 2 — Structural Readability (20 points)

Can a machine find where your answer starts and stops?

These engines use your heading hierarchy to locate and bound answers. Break the structure and you break extraction.

One H1 per page, describing the page. Questions as H2 or H3. Widget labels as spans, not headings. The most common failure I see is a genuinely useful FAQ — real guest questions with real answers, the most citable content on the whole domain — marked up at the least important heading level available, while a booking widget occupies the H1s.

Add tables where content is genuinely comparative. Four neighborhoods with columns for vibe, walkability and typical price is dramatically more extractable than four paragraphs saying the same thing. These engines lift tables cleanly.

Lever 3 — Multi-Modal Content (15 points)

Does anything but text exist?

In Ahrefs' December 2025 study of 75,000 brands, mentions on YouTube correlated with AI visibility more strongly than any other factor measured. Correlation is not proof, but vacation rental companies are sitting on the best raw material in any industry and mostly not using it.

You already have a photo library. You already film property walkthroughs. Publishing those to YouTube with real titles and real descriptions builds the mention footprint conversational engines check. A YouTube icon in your footer with no channel behind it is worse than nothing.

While you are in there, fix your image alt text. Repeating your target keyword verbatim across every image on a page is a 2012 tactic that Google's own image guidelines say can get your site treated as spam. Describe the actual image: "Paddleboarders on the river through the old mill district," not your keyword three times running.

Lever 4 — Authority and Brand Signals (20 points)

Does the world outside your website know you exist?

This is the lever that unlocks the conversational engines, and it lives almost entirely off your own site. In the same Ahrefs study, brand mentions correlated far more strongly with AI visibility than backlinks did. If you have spent years thinking about link building, sit with that for a second. What matters most is not who links to you. It is who talks about you.

Practically:

Get listed where AI looks for lodging. When a guest asks where to stay, AI leans on the people whose job is answering that question: your convention and visitors bureau, your state tourism site, your county's visitor guide, the chamber's lodging page. In one mountain town we audited in November 2026, the visitor bureau's website was the most-used source in ChatGPT's answers to guest questions, 52 uses across 54 answers. In a Midwest market, ChatGPT answered "where is the best place to stay" straight from the county visitor guide and a merchants' association lodging page, and our client was on neither. Most of these listings are free or come with a membership. Get on every one in your market, under the name guests know you by.

Get on the lists ChatGPT reads. In our October 2026 test, managers listed on three or more manager-ranking sites were named in 63% of "best companies" answers; managers on two or fewer, in 23%. That is a correlation, not proof, since stronger companies land on more lists. Look yourself up on each and correct what is wrong.

Fix your Google Business Profile. In the same test, 93% of Google AI Mode's answers about the best rental companies in a market showed Google business cards, and most of the sources in those answers were those cards. Check your category, your reviews and your website link.

Get a Wikidata entry. Covered above. A cheap bet if independent sources describe you.

Participate in community discussion honestly, for Google's sake more than ChatGPT's. In our October 2026 test, social and forum links showed up in 57% of Google's AI Overviews and in 2 of 1,260 ChatGPT answers. The "where should we stay in [your market]" question gets asked in your local subreddit constantly. Answer as a knowledgeable local operator, disclose who you are, follow the subreddit's self-promotion rules, be useful. One genuinely helpful answer a week compounds. Spam does not.

Publish original data. This is the big one, and it is available to you in a way it is not to most businesses. You are sitting on occupancy, average daily rate, booking windows and seasonal demand curves for your market. Nobody else has your data. A quarterly market report makes you a primary source rather than a commentator — it gives every platform, and every journalist, something to cite that nobody else has, and press pickup builds the mention footprint further. If you do one thing from this document beyond the technical fixes, do this.

State your own facts on your own site, and make every copy agree. Founding year, founders, markets served, property count, average rating, the exact company name. I keep finding companies whose founding story exists on six third-party sites and nowhere on their own domain. The opposite problem is just as common: in one November 2026 audit, the company's own pages implied three different founding years, a business directory gave a fourth, and ChatGPT repeated the directory's. Pick one version of each fact, put it on your About page and in your schema, then fix the directories, because AI believes whichever copy it finds.

Lever 5 — Technical Accessibility (20 points)

Can the machines physically get to your content?

Four checks, in descending order of how often I find them broken.

Schema markup. Structured data is how you assert facts in a form machines read without interpretation. Priority order: `Organization` site-wide — including a `sameAs` list pointing at every profile you own, which is how you tell engines which profiles are yours — then `LocalBusiness` per location, `Article` on posts with a named author, `FAQPage` on your FAQ, and `VacationRental` on property pages. Google publishes a structured data type built specifically for short-term rental listings. It requires an `Accommodation` block with occupancy, a stable identifier, exact latitude and longitude, and at least eight photos including a bedroom, a bathroom and a common area; beds and amenities are recommended. Google's own listing feature for it is limited to approved partners, so add the markup because every engine that reads schema can use it, not for a guaranteed Google placement.

Test your actual output. An installed plugin proves nothing. And keep it in proportion: Google says no special markup is required for its AI features. Do this because it states your facts unambiguously, not because it unlocks a feature.

AI crawler access. Confirm you are not blocking the retrieval crawlers, in robots.txt, at your CDN *and* at your host. The names have split into three categories, and this is where the old advice has gone stale:

PURPOSE	CRAWLERS	WHAT TO DO
Retrieval — these decide whether you appear in an answer	`OAI-SearchBot`, `Claude-SearchBot`, `PerplexityBot`	**Allow.** Blocking these is how you disappear from the answer.
User-initiated — fetch a page when a person asks about it	`ChatGPT-User`, `Claude-User`, `Perplexity-User`	**Allow.** OpenAI says robots.txt rules may not apply to these, so when they are blocked it is usually at your CDN.
Model training — these collect content for training corpora	`GPTBot`, `ClaudeBot`, `CCBot`	Your call. Blocking them does not remove you from AI search answers.
Both, and its own special case	`Google-Extended`	Read the note below before you decide.

Anthropic and OpenAI both now run separate retrieval and training crawlers, and blocking one does not block the other. If the training question genuinely bothers you, disallow the training bots and leave the retrieval bots alone. What you must not do is block retrieval out of caution. Wanting ChatGPT to recommend you while blocking its retrieval crawler is like cooking beautiful food and locking the front door.

`Google-Extended` gets confused constantly, including by people selling advice about it, so be precise here. It is **not** a crawler — it is a permission token, and it does **not** control Google Search or AI Overviews. Those run on Googlebot, and blocking `Google-Extended` will not remove you from either one. What it does control is whether your content can be used to train Gemini and to ground Gemini's answers at prompt time. So blocking it costs you Gemini visibility and buys you nothing in Search. If you want to be in Gemini's answers, allow it.

Google's actual switch for AI Overviews and AI Mode is a Search Console setting: Settings, then Search generative AI. It defaults to Include. Confirm nobody has flipped it.

One CDN check: Cloudflare's AI crawler controls changed more than once in 2026, and at least one blocking option for the training category also stops Googlebot, Bingbot and Applebot, because they crawl for both search and training. Whatever you set there, confirm afterward that Googlebot and Bingbot can still reach your pages.

There is a subtler version of this failure that is much harder to spot: a crawler reaches your site, starts, and gives up because your server is too slow. Nothing is blocked, robots.txt is clean, the bot simply never finishes. If your pages are slow, that is no longer only a guest experience problem.

Server-side rendering. If your content only appears after JavaScript runs, assume it does not exist. Load a page with JavaScript disabled. What you see is roughly what a crawler sees.

On a vacation rental site the usual culprit is the booking engine. In both client audits we ran in November 2026, the property search grid loaded its homes with JavaScript, so a crawler saw an empty page, and one site with more than a hundred homes had none of them in its sitemap. Make sure every property page is in your sitemap and linked from a plain list. Turn the filters guests actually use (hot tub, pet-friendly, waterfront, sleeps twelve) into real pages with their own addresses and a list of homes that loads without JavaScript, because those filters are the questions people ask AI. Then check your reviews: a review widget that loads in a frame or a script shows a crawler the star count and none of the words.

llms.txt. A plain text file at your root telling language models what your site is and which pages matter. Here is the honest version, because this file has become the clearest example of how this field gets oversold.

Google states plainly that Google Search ignores it, and that adding one neither helps nor harms you there. No major AI platform has published documentation saying its crawlers use yours to decide what to cite. A few agent tools are starting to look for it. In our own October 2026 test of 132 managers, the ones with an llms.txt file showed no meaningful lift in AI answers.

Where I land: it is a ten-minute job with no downside, so add one. Add it knowing exactly what it is — a cheap bet on something that might matter later, not a lever that works today.

That distinction is worth more to you than the file is. A checklist full of items presented as though they are all doing equal work is precisely how operators end up paying for services that accomplish nothing. Know which of your tactics are proven and which are speculative, and never let anyone blur the two.

And if you already have one, go read it. Some SEO plugins auto-generate these, and an auto-generated llms.txt is worse than none — it will confidently tell every language model that your most important pages are an internal test page and some news items from 2020.

What changes by weight class

Everything above applies to you regardless of size. What changes is sequence.

WEIGHT CLASS	WHERE TO START
Featherweight	Fix your headings — one H1 per page. Get basic `Organization` schema live. Confirm you are not blocking retrieval crawlers. Look yourself up on your visitor bureau's lodging page and the manager-ranking sites, and claim or fix what is there. Create a Wikidata item if independent sources describe you. Add a real byline. A weekend and no money.
Lightweight	Add a direct answer under a question heading at the top of your ten best pages. Convert your FAQ to proper heading levels and add `FAQPage` schema. Start answering questions in your local subreddit. Write real alt text. Publish one season-in-review post a year with your own occupancy and rate numbers.
Middleweight	Build the guide library out to cover every market you operate in — this is your citation engine, and thin coverage in a market where you own properties is money left on the table. Add comparison tables. Turn the search filters guests use most into real pages. Activate YouTube with walkthroughs you are already filming.
Heavyweight	Make your market data quarterly and break it out by market. You have the volume to make it credible, and your competitors do not. Stop competing for citations and become the thing that gets cited.
Super Heavyweight	Build the full entity layer: notability, press, a podcast, a data report journalists reference by name. At this point you are not optimizing for AI search, you are becoming part of the reference material.

What not to do

Do not generate your content with AI. The one thing a language model cannot fake is first-hand experience, which is exactly why engines increasingly weight it. Content that reads as machine-written competes directly against the thing that wrote it, and loses. Your advantage is that you have actually been to the river, met the neighbors, and cleaned the hot tub. Using AI to shape material you already have is a different thing; the book's appendix shows how.

Do not chase every platform. Google's surfaces reward the SEO you are already doing. The conversational engines reward entity work. Pick based on where your guests actually are.

Do not buy GEO services without asking for evidence. This field is young enough that nobody has five years of experience in it, which means anyone claiming deep expertise is telling you something that cannot be true. No credentialing, no settled body of proof — a combination that reliably attracts people who are comfortable making things up. The test: when someone tells you a tactic works, ask what they measured, over what period, and what the result was. If the answer is that it is a best practice, or that everyone is doing it, you have learned what you needed to know. Apply that to me as readily as to anyone else.

Do not treat this as a separate discipline. This is SEO with a wider surface area. Anyone selling you an LLM audit as a distinct product from an SEO audit is probably charging you twice for one job.

Do not wait for it to settle. It will not. There is no stable end state coming where you learn the rules once. What compounds is the underlying work — reviews, real expertise, clean structure, being genuinely known — and every bit of that would be worth doing even if all of these tools vanished tomorrow.

PART 3

Worked example

One example, all the way through.

Fictional company, used to show what a completed scorecard looks like.

Ridgeline Cabin Co. — 40 cabins in one mountain valley. Strong local reputation, ranks well for its main market terms, has never thought about any of this.

LEVER	SCORE	WHY
Citability	**6 / 25**	Destination pages built in a page builder, text split across 40-plus small boxes. No answer blocks anywhere. All blog posts credited to "The Ridgeline Team." No sources cited.
Structural Readability	**4 / 20**	Booking widget renders six H1s on every page. The FAQ — 18 genuinely good guest questions — is marked up at H6.
Multi-Modal	**3 / 15**	Large photo library, all on-site only. Films walkthroughs for listings, publishes them nowhere. YouTube icon in the footer links to a channel that does not exist.
Authority & Brand	**5 / 20**	Not on the visitor bureau's lodging page or any manager-ranking site. No Wikidata entry. No community participation. Founding story lives on a directory site but not on their own About page. Sitting on six years of occupancy and rate data, has published none of it.
Technical	**14 / 20**	Server-side rendered and fast — genuinely good. Retrieval crawlers all allowed. But schema plugin installed and emitting nothing on every page tested, and an auto-generated llms.txt listing a staging page as their most important URL.
Total	**32 / 100**	

Baseline measurement, dated: 15 prompts, run 10 times each across ChatGPT, Perplexity and Gemini. Company named in 2 of 450 responses (150 per tool). Both times in Gemini, never in ChatGPT or Perplexity.

What they should do first, in order, and why: the schema fix and the H1 fix are template-level changes that take a developer a morning and lift two levers at once. Getting onto the visitor bureau's lodging page and the ranking sites costs an afternoon. Wikidata costs an hour. Uploading the walkthrough videos they have already filmed costs an afternoon. Only then does it make sense to start rewriting page content for citability, which is the biggest lever but also the slowest work. The market report is the highest-ceiling item on the list and should be scheduled for next quarter rather than attempted now.

Note what is *not* on that list: nothing about buying a tool, and nothing about llms.txt beyond deleting the auto-generated one that is actively misleading.

PART 4

Your turn

Your worksheet. Print it and fill it in.

Your prompt list

Write ten to twenty questions a real guest or homeowner would type. Full sentences. Include the details a real person volunteers — where they live, who is coming, what they need.

#	PROMPT	AUDIENCE (GUEST / OWNER)
1		
2		
3		
4		
5		
6		
7		
8		
9		
10		

If you struggled to reach ten, stop and deal with that before anything else. It is a positioning problem, not a technical one.

Your baseline

Date measured: _____

TOOL	PROMPTS RUN	TIMES COMPANY NAMED	RATE
ChatGPT			
Perplexity			
Gemini			
Claude			

Next re-measurement date: _____

Your scorecard

LEVER	MAX	YOUR SCORE	BIGGEST SINGLE GAP
Citability	25		
Structural Readability	20		
Multi-Modal Content	15		
Authority & Brand Signals	20		
Technical Accessibility	20		
Total	**100**		

Your first three fixes

Pick from your lowest-scoring levers, but weight toward things that are template-level or free.

#	FIX	OWNER	DUE	DONE
1				
2				
3				

PART 5

How you'll know it worked

What to check, and when.

Within a week — the mechanical checks either pass or they don't. One H1 per page. Schema validating. Retrieval crawlers allowed. JavaScript-disabled page still shows your content. These are binary, so verify them directly rather than waiting on any downstream effect.

Within a month — your answer blocks should be getting picked up somewhere. Watch for your own wording appearing in an AI answer even when you are not linked. That is the earliest real signal, and it usually shows up before traffic does.

Within a quarter — re-run your prompt set and compare the rate against your dated baseline. This is the number that matters. If it has not moved but your score has improved substantially, your work is probably real and the systems have not caught up. Keep measuring before you change course.

What not to read as failure: flat traffic from AI tools. The traffic is small and will stay small for a while. Mentions are the thing you are building, and mentions arrive in your analytics disguised as direct traffic and branded search. GA4's AI Assistant channel shows only the visits that arrive with a referrer, so treat it as a floor. If branded search volume is climbing and you cannot explain why, that may well be this working.

A caution on attribution: you will be tempted to credit or blame this work for revenue changes it did not cause. Almost everything in this document also improves ordinary search performance, which means clean attribution between the two is not really available to you. That is fine — it is also the reason the downside risk here is close to zero.

PART 6

Where this lives in the book

Where to read more in the book.

Chapter: *The Ground Moved Under You — AI Search, GEO, and Getting Found When Nobody Types Into Google.*

Read it before the weight-class chapters, which is where it sits in the book and for good reason: it changes how you should read everything after it. The chapter covers the reasoning behind the Five Levers, why travel answers carry more citations in ChatGPT than any other industry measured, and the measurement problem in more depth than a checklist can.

One instruction from that chapter worth repeating here, because it determines whether any of this is a good use of your time: **optimize for Google first.** Get to solid shape on classical SEO before spending heavily here. The overlap is enormous and the fundamentals are far better understood. This document is not an alternative to the search marketing chapters. It is what you layer on once those are working.

BUILDUP BOOKINGS

Want help putting this to work?

We've worked with more than 300 vacation rental companies on direct bookings, homeowner acquisition, SEO, paid media and AI search. If you'd rather not do this alone, book a free call at **buildupbookings.com**.

All eleven book bonuses: **buildupbookings.com/book-bonuses**

© 2027 BuildUp Bookings. Use it freely in your own business. Ridgeline Cabin Co. is fictional. Platform rules, fees and features change; the facts in this guide were checked in November 2026.