



A CHOOSE-YOUR-OWN-ADVENTURE WORKSHEET

AI Search Readiness Checklist

Somewhere out there, an AI assistant is about to get asked a question your business could answer. Five chapters decide whether it cites you or someone else. Turn the page to begin.

BEFORE YOU OPTIMIZE

Your AI Search Baseline

This is a story with five chapters and three endings — and it's only the beginning. Consider it your baseline: the same foundational checks **JHF Research** runs on every site before recommending anything else. Clear these five chapters and you're ready for the work that actually moves the needle — content strategy, schema architecture, citation building. Check off what's already true, note what isn't, and when you reach the end you'll know exactly which chapter to open next. The checklist ends on the last page; the journey doesn't.

"It's dangerous to go alone. Take this."

— CONSIDER THIS WORKSHEET YOUR SWORD



CHAPTER ONE — THE GATE

Technical Foundations

Every adventure starts with a locked door. If AI crawlers can't get in, nothing past this page matters — turn back and start here.

Why it matters: AI assistants can only cite what they can reach. If a crawler is blocked or your content only renders after JavaScript runs, you're excluded before your content is ever judged on quality.

-
- ☐ robots.txt explicitly allows AI crawlers (GPTBot, ClaudeBot, PerplexityBot, Google-Extended, etc.)
 - ☐ An llms.txt file exists at the root and lists your key pages
 - ☐ XML sitemap is current and submitted to Search Console
 - ☐ Core content renders without requiring JavaScript execution
 - ☐ Pages load fast enough that crawlers don't time out

NOTES





CHAPTER TWO — THE TRANSLATOR

Structured Data & Entity Signals

You can be standing in the room and still go unheard. Schema is how machines translate your site into facts they're willing to quote.

Why it matters: structured data gives AI systems a direct, unambiguous read on what your page means instead of forcing them to guess. Sites without it are easier to misread, mis-cite, or skip entirely.

-
- ☐ Organization schema (name, logo, sameAs links) on the homepage
 - ☐ FAQPage or HowTo schema on pages that answer common questions
 - ☐ BreadcrumbList schema reflecting your site hierarchy
 - ☐ Author/Person schema on articles and guides
 - ☐ Schema validates cleanly in Google's Rich Results Test

NOTES



CHAPTER THREE — THE ORACLE'S TEST

Content Built for Direct Answers

An AI assistant is an impatient oracle — it wants the answer, not the journey to it. Pages that bury the point get skipped for ones that don't.

Why it matters: AI assistants extract answers, they don't read your whole page top to bottom. Content that buries the point rarely gets pulled into a citation, even when the information itself is accurate.

-
- ☐ Each page answers its core question in the first 1–2 sentences
 - ☐ Headers are phrased as the questions readers actually ask
 - ☐ Claims are backed by specific data, examples, or sources
 - ☐ Content shows a visible last-updated or publish date
 - ☐ One page = one clear topic, not one answer buried in a catch-all

NOTES



CHAPTER FOUR — PROOF OF WORTHINESS

Authority & Citation Signals

Before an AI vouches for you, it looks for evidence that someone else already did — a second party willing to say you're the real thing.

Why it matters: AI models weigh third-party validation heavily when deciding what to trust. A site that only talks about itself is a weaker source than one other credible sites vouch for.

-
- ☐ About/Team pages establish real expertise and credentials
 - ☐ Site is cited or linked from other reputable sources in your space
 - ☐ Business name, address, and details are consistent across the web
 - ☐ Reviews or testimonials are visible and specific

NOTES





Tracking What's Working

Every good adventurer keeps a map of where they've been — otherwise you'll never know which chapters actually paid off.

Why it matters: most teams can quote their Google ranking down to the position but have no idea whether ChatGPT or Perplexity ever mention their brand. You can't improve what you don't measure.

-
- ☐ Referral traffic from AI assistants (chatgpt.com, perplexity.ai, etc.) is tracked separately in analytics
 - ☐ You periodically check whether AI assistants cite or mention your brand
 - ☐ You review which of your pages get cited, and why

NOTES

CHOOSE YOUR ENDING

Count your checked boxes, then find your ending below.

Out of 22 total, how many did you check? Your total decides which ending you're living in right now — and what the next chapter should be.

THE ROGUE

0–7

Effective in the Shadows

You're getting real results on instinct and scrappy effort alone.

That's not nothing — but you're invisible to the systems that decide who gets cited, working off the books instead of in the open. **Next chapter:** it's time to level up. Start with Chapter One and Two.

THE ADVENTURER

8–15

The Well-Trodden Path

You're on the map, but not the highlighted route. Real signals are in place; AI assistants can probably find you and sometimes trust you.

Next chapter: go back to whichever section had the most unchecked boxes — that's your fastest path to the next ending.

THE WIZARD

16–22

The Main Quest, Complete

You're ahead of most competitors. The foundations, the structure, and the proof are all there. **Next chapter:** this is a maintenance story now — keep content current and keep watching Chapter Five to see what's actually getting cited.