



APPWT WEB & AI SOLUTIONS

MICHIGAN, SINCE 1997

The AppWT AI Visibility Playbook

The System for Getting Your Business Cited by AI

Written by Anthony “Tony” Paris

AppWT Web & AI Solutions (AppWT LLC, a family company)



The AppWT AI Visibility Playbook

The System for Getting Your Business Cited by AI

Written by Anthony “Tony” Paris

Published by AppWT Web & AI Solutions

AppWT LLC, a family company

The AppWT AI Visibility Playbook

The System for Getting Your Business Cited by AI

Written by Anthony "Tony" Paris

AppWT Web & AI Solutions (AppWT LLC, a family company)

Second Edition, August 2026

© 1997-2026 AppWT Web & AI Solutions (AppWT LLC). All rights reserved.

Published in Livonia, Michigan, United States of America.

A NOTE ON SPELLING

This manual is read in the United States, Canada, the United Kingdom and Australia. Where a word is spelled differently in different countries, both spellings appear the first time the word is used, like color/colour. After that, one spelling is used so the page stays easy to read.

HOW TO USE THE TECHNICAL NOTES

Smaller italic notes marked *For your tech person* hold the deeper detail: exact settings, file names and commands. You can skip every one of them and the manual still makes complete sense. Hand them to whoever looks after your website.

IMPORTANT

This manual describes security practices that AppWT uses and recommends. It is written as general guidance for business owners. It is not legal advice, and it is not a guarantee that any system cannot be attacked. Security reduces risk. No honest provider will promise to remove risk entirely.

Every incident described in this manual comes from work AppWT actually performed. Client identities are deliberately withheld, because a client's security incident is not ours to publish. Where the business involved is AppWT itself, we name ourselves.

A Note From Tony

For twenty-nine years the question was where you ranked. That question is quietly becoming the wrong one.

When somebody asks an assistant which company to call, they do not get ten links. They get an answer, assembled out of a handful of sources, and a couple of those sources get named. You are either in that answer or you are invisible, and there is no page two to be on.

This book is the system we use to get businesses into those answers. Not the theory of it. The files, the patterns, the page templates and the tracking sheet.

I want to be straight with you about two things, because a lot of what is written on this subject is not.

The first is that nobody controls what an AI engine says. Anyone selling you a guarantee is selling you something they cannot deliver. What you can control is whether the engine can reach your pages, understand what you are, and find a clean sentence worth quoting. That is most of the battle, and almost nobody has done it.

The second is that a good deal of this is genuinely new, and some of it will change. Where a practice is settled, I say so. Where it is emerging and we are betting on it, I say that instead. Chapter 17 is entirely about the things that are coming, and it is labelled as such so you can tell the difference.

Work through it in order the first time. The files in Part Two take an afternoon and they are the difference between being readable and being skipped. Part Three looks like the hardest section and is actually the most mechanical, because you are copying patterns. Part Four is where the real writing happens.

Then keep the tracking sheet in Chapter 15 open. Visibility is not a thing you finish. It is a thing you watch.

Under-promise, over-deliver. That has been the rule here since 1997, and it is why this book tells you what will not work as carefully as it tells you what will.

Anthony “Tony” Paris

Founder, AppWT Web & AI Solutions · Livonia, Michigan

How to Use This Manual

This is a working book. It is meant to be open beside you while you change things, not read once and shelved.

The five parts

Part One explains how an AI engine chooses its sources. Read it once, straight through. It is short and it makes everything after it make sense.

Part Two is four files. If you do nothing else in this book, do Part Two. It is the difference between an engine being able to read your site and quietly skipping it.

Part Three is structured data. It looks technical and it is mostly copying. Each chapter gives you the pattern to paste and tells you what to change.

Part Four is the writing. This is where most businesses lose, because the page is fine for a human and useless to a machine.

Part Five is measurement, a thirty-day plan, what is coming next, and a list of advice to stop following.

The boxes, and what each one means

DO THIS TODAY

A gold box is one action you can finish now. If you only ever act on the gold boxes, you will still be far safer than when you started.

FROM OUR FILES

A dark box is a real incident from AppWT work. Nothing in these boxes is invented. Where a detail was never verified, we say so rather than fill the gap.

FOR YOUR TECH PERSON

A smaller italic note like this one carries the professional detail: exact file names, headers, settings and commands. Skip every one and the manual still reads completely. Forward them to your web person.

About the code

Patterns in this book are written to be copied and edited. Every one has been kept short enough to read on paper. Where you need to change something, the text says exactly which part.

If you are handing the technical work to somebody else, the smaller italic notes and the code are the parts to send them. The main text is yours.

A word about promises

No book can make an AI engine cite you, and this one does not claim to. What it does is remove every reason for an engine to pass you over. That is a real advantage and it is one almost nobody has bothered to take.

If you only have an afternoon

Do Part Two and Chapter 12. Four files and one page template. That combination has moved more sites into AI answers than anything else in this book.

Contents

PART ONE *How AI Decides What to Cite*

1	From Ranking to Being Cited	2
2	The Engines That Matter, and How Each One Finds You	9
3	What a Citation Actually Requires	16

PART TWO *The Files That Make You Readable*

4	robots.txt and the AI-Crawler Allowlist	23
5	llms.txt and llms-full.txt	30
6	The Rest of the File Set	37

PART THREE *Telling Machines Who You Are*

7	Structured Data, Part One: Who You Are	45
8	Structured Data, Part Two: The Page Itself	53
9	Structured Data, Part Three: Answers and Actions	60
10	Speakable and Voice	68
11	Person, Authorship and E-E-A-T	75

PART FOUR *Pages Built to Be Quoted*

12	The Answer Page Template	82
13	Writing So a Machine Can Quote You	91
14	Question Research Without Guesswork	98

PART FIVE *Measuring It and Keeping It*

15	The Eight-Engine Citation Tracking Sheet	107
16	The Thirty-Day Plan	117
17	The Frontier: From Cited By AI to Used By AI	131
18	What Not To Do	139
	One-Page Checklist	147
	Glossary	149
	About AppWT	151



PART ONE

How AI Decides What to Cite

What actually changed when people started asking assistants instead of scanning search results, and what an engine needs before it will quote you.



From Ranking to Being Cited

The shift from a place on a results page to a place inside the answer

People ask an assistant now instead of scanning ten blue links, and the goal moved with them. This chapter explains what GEO, AIVO, LLMO, AEO, VSO and SEO mean, gives you three tests you can run today, and sets up the four working parts of this playbook: files, structured data, answer pages, and measurement.

The goal of website visibility has moved. It used to be a position on a page of results. It is now being the source an AI assistant builds its answer from. A person who asks a chatbot a question often never sees a list of ten links, so a rank on that list cannot help them find you.

Being cited means an assistant used your page to write its answer. Sometimes your business name appears in the sentence. Sometimes a small source link sits beside the paragraph. Either way, the assistant reached your page, trusted what it found, and handed your information to a real person.

The scoreboard changed. The work did not vanish.

WHAT YOU WATCH	SEARCH ERA ANSWER	ASSISTANT ERA ANSWER
The win	Rank number three for a keyword	Named as a source inside the answer
The unit that competes	The whole page	One passage that answers one question
Who reads it first	A person scanning links	A model pulling out facts

WHAT YOU WATCH	SEARCH ERA ANSWER	ASSISTANT ERA ANSWER
What hurts you most	Sitting on page two	Being unreachable, unreadable, or undated
The proof	A rank tracker	A logged check of real prompts across engines

Do not throw away search engine optimization/optimisation. To be cited in a Google AI Overview, a page normally has to rank in the top ten for that query first. Microsoft's assistant works from the Bing index, so strong Bing performance and strong assistant performance are the same work. Classic SEO is the ticket into the room where citation happens.

Google AI Mode goes a step further. It splits one question into several smaller ones, then builds a single answer from all of them. Deep coverage of a topic helps you here. So do passages that make sense on their own, because more of them get pulled into the mix.

Six labels, one job

LABEL	FULL NAME	WHAT IT MEANS IN ONE LINE	WHERE THE RESULT SHOWS UP
GEO	Generative Engine Optimization	Writing so a chat assistant can quote you correctly	Grok, ChatGPT, Claude, Gemini and Perplexity answers
AIVO	AI Visibility Optimization	Making sure AI systems can find and identify your business at all	Brand mentions, knowledge panels, directory data
LLMO	Large Language Model Optimization	Publishing accurate, lasting content a model learns from over time	What an assistant says about your field with no live search
AEO	Answer Engine Optimization	Formatting a direct answer a machine can lift whole	Featured snippets, People Also Ask, AI Overviews
VSO	Voice Search Optimization	Writing answers that still make sense read aloud	Siri, Alexa, Google Assistant and Bixby replies

LABEL	FULL NAME	WHAT IT MEANS IN ONE LINE	WHERE THE RESULT SHOWS UP
SEO	Search Engine Optimization	Earning a place in the index the others draw from	Blue links, and the pool AI answers pull from

FOR YOUR TECH PERSON

Treat the six labels as one discipline seen from six angles, not six separate budgets. Every tactic in this book has passed a durability filter first. If a practice will not still matter in 2027 or 2028, it is not in here. That filter is why you will not find keyword density targets, AMP, or the meta keywords tag anywhere in these pages. It is also why the speed numbers here are the current Core Web Vitals: Largest Contentful Paint under 2.5 seconds, Interaction to Next Paint under 200 milliseconds, and Cumulative Layout Shift under 0.1. Interaction to Next Paint replaced First Input Delay in March 2024, so ignore any guide still tuning for the old metric.

A rank is a place you occupy. A citation is a sentence somebody else writes with your name in it.

What an assistant needs from your page

1. Access. A crawler is a program that fetches your pages for a search or AI company. It has to be allowed in, and it has to actually receive the page.
2. Plain HTML. The answer must be in the page source. JavaScript is code that builds part of a page inside the visitor's browser after the page arrives. Content painted that way is often invisible to a crawler.
3. An answer near the top. Put the short answer in the first 100 words, then give the detail underneath.
4. One claim per paragraph. Short, factual, standalone paragraphs are the easiest to quote without changing your meaning.
5. Labels a machine can read. Structured data is a small block of machine-readable facts added to your page in a format called JSON-LD. It tells an engine what the page is and who you are.
6. Honest dates. A published date and a last updated date, both accurate. Freshness is a real inclusion signal, and a fake date is worse than no date.

TEST 1. CAN A CRAWLER ACTUALLY GET YOUR PAGE?

Run this in a terminal. Swap in one of your own page URLs.

```
curl -A "GPTBot" -s -o page.html \
```

TEST 1. CAN A CRAWLER ACTUALLY GET YOUR PAGE? (CONTINUED)

```

-w "status=%{http_code}\n" https://example.com/your-page/

# 1. status must be 200. Anything else means the page was refused.
# 2. Check that your answer text is really in the raw HTML:

grep -c "your exact answer sentence" page.html

# A result of 0 means people see the text but crawlers do not.
# 3. Repeat for the assistants you care about:

for UA in "Grok" "GPTBot" "OAI-SearchBot" "ClaudeBot" "PerplexityBot"; do
  echo -n "$UA -> "
  curl -A "$UA" -s -o /dev/null -w "%{http_code}\n" \
    https://example.com/your-page/
done

```

A status code other than 200 usually means something in front of your website refused the request. A user agent is the name a program gives when it asks for a page. A web application firewall is a filter that inspects traffic before it reaches your site. Security suites and firewalls often challenge a user agent they do not know, and AI crawlers get caught in that net. The filter can answer with a challenge screen or an error instead of your page. Ask your host or security vendor to allow the crawlers you want. Then run the same test again and confirm a clean 200.

FOR YOUR TECH PERSON

Two different kinds of visit now arrive at your server. An indexing crawler reads your site to build a library for later. An agent fetches one page right now because a person just asked it to do something. Agentic browsers are already here, including ChatGPT agent and Operator, Google Project Mariner and AI Mode, Perplexity, and Anthropic computer-use. Agents read the accessibility tree, so they need stable semantic HTML with ARIA labels, single-purpose forms with correct labels and autocomplete tokens, and predictable URLs. Keep any challenge screen off the main path and put a JSON endpoint behind each key action where you can. Agent experience is the new mobile-first. Log both kinds of visit and review your allow list on a set schedule, because new AI crawlers keep launching.

The four working parts of this playbook

PART	WHAT IT COVERS	THE FIRST THING YOU MAKE
Files	The small text files at your site root that tell machines who you are and what they may read	robots.txt, sitemap.xml, llms.txt and llms-full.txt

PART	WHAT IT COVERS	THE FIRST THING YOU MAKE
Structured data	JSON-LD blocks naming your business, pages, services, questions and media	Organization and WebSite schema on the home page
Answer pages	Pages built around real questions, answer first, one claim per paragraph	One FAQ block holding five real customer questions
Measurement	A repeatable check of which engines cite you, for which prompts	A fixed prompt list and a monthly tracking sheet

TEST 2. A STARTER ROBOTS.TXT THAT LETS THE ASSISTANTS IN

```
# robots.txt - save at https://yourdomain.com/robots.txt
# Grok is first because xAI reads X posts in real time and
# tends to pick up new pages faster than the rest.
```

```
User-agent: Grok
User-agent: GPTBot
User-agent: ChatGPT-User
User-agent: OAI-SearchBot
User-agent: ClaudeBot
User-agent: anthropic-ai
User-agent: PerplexityBot
User-agent: Google-Extended
User-agent: Applebot-Extended
User-agent: Amazonbot
User-agent: CCBot
Allow: /
```

```
User-agent: *
Allow: /
Disallow: /admin/
Disallow: /cart/
```

```
Sitemap: https://yourdomain.com/sitemap.xml
```

HOW APPWT DECIDES A PAGE IS ACTUALLY FINISHED

AppWT does not call a page optimized until twelve checks pass. The list covers exactly one h1 per page, semantic HTML, modern image and video formats with working fallbacks, a complete schema graph, Speakable markup for voice, ARIA and WCAG 2.1 AA accessibility, and real search intent in the title and first 100 words. It also covers no two pages presenting the same information, AI crawler access with llms.txt and llms-full.txt shipped, voice answers of 40 to 60 words, a unique title and description and canonical with the sitemap submitted, and a sitemap never more than three days stale. The most important word in that gate is the last one. Verification happens against the live URL, fetched with a real crawler user agent. Never against a staging copy, and never against an assumption.

How to tell if it is working

The number to watch is citation share, also called AI citation rate. It is the share of your chosen prompts where at least one major assistant cites your site. Tools for this are appearing, including Profound, Peec AI and Otterly.AI. Most of them run your keywords through the assistants and track citation rates over time. Run your own check by hand as well, so you always hold a first-party number. Analyze/analyse the direction of that number across months rather than fixating on its size in any one month.

YOUR FIRST WEEK, IN ORDER

Day 1. Run Test 1 on your five most important pages and write down every status code. Day 2. Publish or repair robots.txt using Test 2, then confirm your sitemap URL returns 200. Day 3. Add Organization and WebSite structured data to the home page and validate it. Day 4. Rewrite the top of your three strongest pages so the answer sits inside the first 100 words. Day 5. Write down twenty questions customers actually ask you, then pick the five you answer best. Day 6. Start the tracking sheet below and run your first check. Day 7. Add nothing new. Re-run Test 1 and confirm nothing broke.

TEST 3. THE CITATION TRACKING SHEET

```
# ai-citation-log.csv - one row per prompt, per engine, per check.
# Run the same prompts on the same day each month.

date,prompt,engine,cited,rank_in_answer,url_cited,notes
2026-08-01,"emergency plumber in Livonia MI",Grok,yes,2,https://example.com/,"cited home not city page"
2026-08-01,"emergency plumber in Livonia MI",ChatGPT,yes,1,https://example.com/livonia/,"named in first line"
2026-08-01,"emergency plumber in Livonia MI",Perplexity,yes,3,https://example.com/livonia/,""
2026-08-01,"emergency plumber in Livonia MI",Gemini,no,,,"no local sources used"
```

TEST 3. THE CITATION TRACKING SHEET (CONTINUED)

2026-08-01,"emergency plumber in Livonia MI",Claude,no,,, "generic answer, no sources"

cited: yes or no, nothing else.

rank_in_answer: 1 means you were the first source named.

Monthly score = rows marked yes divided by total rows.

That percentage is your citation share. Watch which way it moves.

WHAT TO REMEMBER

- A rank is no longer the goal. Being the source an answer is built from is the goal.
- GEO, AIVO, LLMO, AEO and VSO are five angles on one job, and classic SEO is still the ticket into the room.
- Check access before you check anything else. A page a crawler cannot fetch will never be cited, no matter how well it is written.
- Put the answer inside the first 100 words, keep one claim per paragraph, and date the page honestly.
- Measure citation share by hand, monthly, against a fixed prompt list. The direction of the number matters more than its size.



The Engines That Matter, and How Each One Finds You

Chapter 2.

Eight AI engines decide whether an answer names your business. This chapter gives you the order to work them in, the exact user agents to allow, what each engine needs before it will cite you, and the commands that prove the work is real.

Eight AI engines decide whether an answer names your business. Work them in this order: Grok, then ChatGPT, Claude, Gemini, Copilot, Perplexity, DeepSeek and You.com. That order is not about company size. It is about how fast you can see the result of your work. Grok shows it first, so Grok goes first.

Grok Leads. Here Is Why.

Grok is the assistant built by xAI and wired into X, the site formerly called Twitter. Grok reads X's real-time firehose, which is the live stream of public posts. A page linked in a public post can reach Grok faster than any other assistant picks it up. X posts are also a ranking surface of their own, so one post does two jobs. When you optimize/optimize for Grok first, you get the fastest honest proof that a change worked.