

AI Findability Scan

Harborline Cruises

July 28, 2026

Where You Stand

Across ChatGPT, Claude, and Gemini, **Harborline Cruises is the most-named harbor cruise brand in Crescent Bay**. It is recommended by name in **82% of the buyer questions we tested**, and it ranks **first** ahead of every competing operator on all three engines.

This is not a spot check. We asked **40 real buyer questions** about harbor cruises in Crescent Bay, put every one to **ChatGPT, Claude, and Gemini**, and repeated each question across **three separate sessions**. That is **360 scored answers**, measured for one thing: when a buyer asks, does the AI name you. The short answer is yes, consistently, and more often than anyone else. Harborline is not one option in a long list. On most questions it is the brand the AI reaches for first.

How each engine names you

AI engine	Names Harborline by name	Standing vs competitors
ChatGPT	78% of questions	1st
Claude	83% of questions	1st
Gemini	84% of questions	1st
All three (blended)	82% of questions	1st

Share of answers that named Harborline Cruises by name, across three engines and three sessions (360 answers).

Why: The Diagnosis

You own the questions where a buyer already knows what they want, comparisons, reviews, and direct recommendations, at or near **100%**. The one place the standing softens is **open-category discovery**, the broad questions where a buyer names no brand at all.

Not every buyer question is the same. Some name you or a rival directly; some ask for the best; some are wide open. Breaking the 40 questions down by type shows exactly where your standing is locked in and where there is still open ground.

Win rate by buyer-question type

When a buyer asks...	Example question	Names you
To compare two operators	"Harborline vs Cascade Water Tours"	100%
For reviews or reputation	"Is Harborline Cruises any good"	100%
Who they should book	"Who should I book for a harbor cruise"	100%
For the best in the category	"Best harbor cruise in Crescent Bay"	94%
For open-ended ideas	"Fun things to do on the water in Crescent Bay"	70%

Share of answers that named Harborline, by question type, across three engines and three sessions.

The standing is steady, not a fluke. Repeating every question across three sessions, Harborline held its position on 80% of questions on all three engines at once (ChatGPT 88%, Claude 90%, Gemini 92%). The small amount of movement clusters in the same open-ended discovery questions, where the answer set is widest and AI occasionally reaches for someone else.

Who Is Winning, and Why

No competitor comes close. The nearest operator is named in **under 4%** of answers, against your **82%**. Your standing rests on being the brand AI describes by name, not one entry in a list.

We tracked every operator the three engines named across all 360 answers. Four competitors surfaced at all. Together they account for 42 mentions; Harborline accounts for 294.

Who AI names in the harbor cruise category

Operator	Share of AI mentions	Where it shows up
Harborline Cruises (you)	82%	Every question type, all three engines
Cascade Water Tours	3.9%	Open-ended discovery, mostly on ChatGPT
Harbor City Cruises	3.3%	Comparison and dining-cruise questions
Bayline Ferries	2.5%	Transport and route questions
Waterfront Events Co.	1.9%	Private-event questions, one engine only

Share is the percent of 360 answers that named each operator. An answer can name more than one operator, and many name none.

The pattern behind the standing

The reason is consistency of naming. All three engines describe Harborline the same way, an established harbor cruise operator with a clear set of experiences, so it is the safe answer whatever the question. Competitors surface mainly when a question is broad enough that AI lists several options, or when the buyer names the competitor directly. On a question about a specific rival, that rival is naturally the subject, and even there Harborline shows up as an alternative about one time in five.

Moves Worth Making

You are the most-named brand today. The site signals AI leans on to keep choosing you score **49 out of 100**. That is your headroom, and your insurance as competitors invest. The plan below is **eleven changes across 90 days**, ordered by lift for effort.

Visibility and site readiness are two different things. Harborline already wins the answer; the readiness score measures how easy your own site makes it for AI to keep naming you, and to name you on the broad questions where the standing sits at 70% instead of 100%. Today that score is 49 out of 100, strongest on answer readiness and with the most room in trust and authority. Two things are already working in your favor: a comprehensive FAQ section and a clear contact page.

AI-readiness score, by pillar

Pillar	Score	What it measures
Answer readiness	61 / 100	How directly your pages answer real buyer questions
Content structure	42 / 100	Whether AI can parse and lift your content cleanly
Trust & authority	39 / 100	The signals that tell AI your answers are reliable
Overall	49 / 100	Weighted across all three pillars

Readiness is scored from a 50-page crawl of your site against 40 answer-engine checks. It is separate from the visibility numbers above: you can rank first today and still have readiness to build.

The 90-day plan, in priority order

#	Move	Effort	Priority
1	Publish an llms.txt file so AI knows what your site is and which pages matter	Low	High
2	Add structured data (Product, Event, LocalBusiness) to your core pages	Medium	High
3	Rebuild heading hierarchy so each page leads with the answer	Low	High
4	Align name, address, and phone everywhere they appear	Low	Medium

5	Publish one "Harborline vs the alternatives" comparison page	Medium	High
6	Add author and expert attribution to your key pages	Medium	Medium
7	Sharpen entity clarity so AI never confuses you for a look-alike	Medium	Medium
8	Strengthen internal linking across your experience pages	Medium	Medium
9	Earn more third-party authority sources that name you	High	Medium
10	Surface your third-party reviews on your own pages	Medium	Medium
11	Add visible freshness signals (dates) to core content	Low	Low

Each move maps to a specific readiness check that scored below its potential in the crawl. Effort is a rough build estimate; priority weights impact against effort.

First 30 days: the quick wins

Four changes cost little and move the readiness score the most. **An llms.txt file** and **structured data** are the two biggest single levers: neither exists on your site today, and both give AI a clean, machine-readable description of what you offer, which is exactly what it reaches for on open-ended questions. **Fixing heading hierarchy** so each page leads with its answer, and **aligning your name, address, and phone** across the web, are fast trust signals that round out the first month. **Expected effect:** more consistent naming on the broad discovery questions where the standing sits at 70% today.

Days 30 to 60: build the assets AI reaches for

Buyers who compare operators already get Harborline named 100% of the time, but there is nothing on your site for AI to cite when it does. A single **comparison page**, written to lead with the answer, gives AI a source to quote and protects that standing as competitors put money into their own content. Alongside it, **author and expert attribution**, **sharper entity clarity**, and **stronger internal linking** all raise the trust and structure pillars, the two with the most room today.

Days 60 to 90: widen the authority base

The last three moves compound over time. Today one strong outside source carries most of your authority; **earning a few more third-party citations** and **surfacing your reviews on your own pages** give AI more independent proof to lean on. **Visible freshness dates** are the final, low-effort signal that keeps your content reading as current. None of these is urgent, and together they are what keep you the default recommendation a year from now.

What the Full Read Adds

This Scan covers three engines and the shape of your standing. The full **Muse Presence** read goes wider and deeper, and the amount you paid for this Scan applies in full toward it.

The Scan is the fast read. The full Muse Presence engagement adds:

- Full three-session stability scoring across 50 or more buyer questions, not the capped taste shown here.
- A claim-by-claim accuracy check on how AI describes you, against verified facts.
- The community, video, and Google AI answer surfaces this Scan deliberately leaves out.
- The five-dimension framework and your Share of Model Voice headline.
- A 90-day action plan with the full priority stack, and the Implementation Pack to execute it.

Our guarantee: every Scan surfaces at least one move worth making, or it is free. The amount you paid for this Scan applies in full as a credit toward a Muse Presence engagement.