



Make AI Work for You

Why, What, and How we make AI work for you

Customers now arrive through AI. AI recommends what it can read.

The shift is already measurable in travel, the category where AI planning took off first.

AI SENDS MORE CUSTOMERS

+119%

AI-referred traffic to US travel sites, July 2026 vs a year earlier. Up 1,822% since October 2024.

+21%

more engaged

+67%

longer visits

-42%

lower bounce

Adobe Digital Insights, August 2026

AND THEY ACT ON IT

45%

of travelers who used generative AI to plan a trip booked accommodation the AI recommended.

Deloitte, 2025 Travel Industry Outlook (survey of 193 GenAI trip planners)

BUT MOST SITES AREN'T READY

68%

machine-readability of hotel product pages: nearly a third of the content is invisible to AI. Hotel homepages score 71%.

Adobe AI Content Visibility Checker, travel sites, August 2026

If AI can't read it, AI can't recommend it. The rest of this deck is how we make sure it can, and prove it on a fixed test.

SECTION ONE

Why AI visibility matters now, and **why us.**

What is at stake when customers ask AI instead of searching, what it takes to be the answer, and how we compare.

01

What AI sees today

02

What strong AEO requires

03

Us vs the market

04

How you will know it worked

What AI sees on your site, before we start.

[Business name] · [website URL] · site check run [date]

01 · AI VISIBILITY

AI visibility

Do ChatGPT and Gemini name and cite you on 150 customer questions?



Measured in the full audit

150 questions × 2 engines × 5 runs = 1,500 answers, against the competitors you name.

02 · AI CRAWLABILITY

AI crawlability

Can AI crawlers reach, fetch and read your pages? The audit delivers the fix: llms.txt and FAQ schema.

00 / 100

[What we found, in one line]

03 · CONTENT STRUCTURE

Content structure

Are facts in text AI can lift: question headings, FAQs, structured data?

00 / 100

[What we found, in one line]

04 · TECHNICAL & AUTHORITY

Technical & authority

Speed, schema, consistent facts across the web: the signals that say "trust this".

00 / 100

[What we found, in one line]

WHERE TO START

[Weakest area, why it matters for this business, and the first fix we would make]

Three of four measures from our site check, run before this meeting. Each scored separately out of 100; we do not roll them into one number.

An engine will only recommend what it can read, trust, match and measure.

Six requirements. Miss one and the answer goes to a competitor or a third-party listing.



01

Readable facts

Every claim as a plain, checkable statement on a page engines can fetch. Not inside images, PDFs or adjectives.

WHERE MOST BUSINESSES ARE TODAY

Key facts live in brochures and photos. The site says "premium" but not what, when or how much.



02

One truth everywhere

Your site, listings and directories agree. When sources conflict, engines hedge or quote the third party instead of you.

WHERE MOST BUSINESSES ARE TODAY

Third-party listings are stale or wrong, and they win the answer.



03

Answers to real questions

Content that answers what customers actually ask AI, in their words. Engines match questions, not slogans.

WHERE MOST BUSINESSES ARE TODAY

Nobody knows which questions are being asked, so nothing on the site answers them.



04

Reasons to choose you

Specific, verifiable differentiators engines can use to recommend you over the next option.

WHERE MOST BUSINESSES ARE TODAY

Claims are generic. Nothing separates you from the alternatives.



05

Kept current

Facts that stay true. Engines down-weight sources that go stale and keep what keeps changing.

WHERE MOST BUSINESSES ARE TODAY

Updated once at launch, never again.



06

A score, not an opinion

You know how often engines name you, cite you and get you right, on a fixed test you can repeat.

WHERE MOST BUSINESSES ARE TODAY

No one has ever checked what AI says about the business.

Most AEO offers cover one requirement. Strong AEO needs all six.

What businesses are typically sold today, what each leaves out, and where the Tabi Life stack sits against it.

WHAT THE MARKET SELLS



AI-content consultants

Write "AI-optimized" pages and FAQs for the client's site.

WHAT IS MISSING

No measurement of what AI actually says. Wrong facts on third-party sites stay wrong. Stops when the content budget stops.



Monitoring dashboards

Track how often an engine mentions the brand, with a score.

WHAT IS MISSING

A number, not a fix. No analyst review, no root cause, nothing published for AI to cite. The client is left to act alone.



SEO agencies adding AEO

Schema markup, llms.txt and technical checks bolted onto SEO.

WHAT IS MISSING

Makes the site readable, not accurate or distinctive. No competitor benchmark, no question set, no growth loop.

WHY WE ARE AHEAD

The market sells one piece. We run the whole loop, and prove it on the same test.

Diagnose which questions you lose and which facts are wrong. Fix them at the source. Publish one verified profile AI can cite. Keep finding the next question through the Learner, GEM with AI and AI Chat. Re-run the identical audit to show the gain.

[DIAGNOSE](#)
[FIX AT THE SOURCE](#)
[PUBLISH](#)
[GROW](#)
[PROVE](#)

WHAT STRONG AEO REQUIRES	CONTENT CONSULTANT	MONITORING TOOL	SEO + AEO AGENCY	TABI LIFE	HOW
01 Readable facts					22 technical checks; llms.txt and FAQ schema delivered ready to install; one plain sentence per fact on your AI Profile
02 One truth everywhere					Every mention read by an analyst; each error traced to its source and fixed there; third-party listings checked against your site
03 Answers to real questions					150 questions in your customers' words; 12 articles a year written from the audit briefs and real unanswered questions ARTICLES FROM NOVEMBER
04 Reasons to choose you					Benchmarked against your named competitors on the same questions; only-here facts verified and published on your AI Profile
05 Kept current					AI Profile updated after one approval, no website release; new questions every week from the Learner, GEM with AI and AI Chat
06 A score, not an opinion					150 frozen questions × 2 engines × 5 runs vs named competitors; identical re-audit; only changes outside the 95% confidence interval count

covered partial not offered - Typical scope of each vendor type; individual providers vary.

A fix found today **shouldn't wait for a website release.**

AI, customer questions and competitors keep changing. How fast you can update decides how much of the work pays off.

THE USUAL ROUTE · CHANGES THROUGH YOUR WEBSITE

Gap found



Brief your web vendor



Wait for a quote



Internal approval



Build



Review



Publish

Every update becomes a project, so updates slow down, then stop.

WITH TABI LIFE · CHANGES THROUGH YOUR AI PROFILE

Gap found



We confirm the fact with you



You approve it in the AI Learner



AI Profile updated



Re-measured at the next audit

One approval from your team. We do the rest, and measure the result.



No vendor brief, quote or release

Facts go live without touching your site.



A few questions a week

Your team answers and approves; no writing.




Your website still matters

We still give you the fixes for it; AI doesn't have to wait.

No one can trace revenue to a single AI answer. So we build a chain of evidence.

Five links, each measured on its own. Together they show what changed, why, and what it is worth to you.




THE QUARTERLY RE-AUDIT SCORECARD			
MEASURE	BEFORE	AFTER	GOAL
AI mentions	baseline	Q1 re-audit	↑
Citations of your own site	baseline	Q1 re-audit	↑
Share of voice vs competitors	baseline	Q1 re-audit	↑
Questions where you never appear	baseline	Q1 re-audit	↓
Answers with an error about you	baseline	Q1 re-audit	↓
Questions nobody answers yet	baseline	Q1 re-audit	↓

Only changes outside the 95% confidence interval count as change.

WHAT WE WILL NOT DO

Put an invented revenue figure on a slide.

Most AI answers pass no referrer, so anyone claiming exact attribution is estimating. We show you every link in the chain, measured against a written baseline and an identical re-audit, and you judge the value with your own economics.



SECTION TWO

What we do: **deliver** every piece strong AEO requires.

Six requirements, each with an owner in the stack and on the team, run as one loop and proven on the same test.

01

Readable facts

02

One truth everywhere

03

Answers to real questions

04

Reasons to choose you

05


Kept current

06

A score, not an opinion

Your AI voice, broken into the parts we measure and move.

One number is not enough. Each part has its own measurement, its own lever, and its own place in the dashboard.



Visibility, accuracy, distinctiveness and reach say where the voice is. Missing voice says where it needs to be next. Growth proves the loop works.

More tags answer more questions. More answered questions give you a bigger **voice**.

A tag is one verified fact about you, in a form AI can use. Our job is to grow that database and widen the reach of your voice with it.



How the tag database grows

Questions customers ask the AI Chat on your site that it cannot yet answer. Baseline facts from Site Intake. Gaps the audit exposes. What customers ask through GEM with AI. The questions the Learner puts to you each week. Every answer you approve becomes a tag.



Why it compounds

One tag answers many questions. Every question you can answer is a moment AI can cite you instead of someone else. Each re-audit shows the questions still unanswered, which become the next tags.



What that means for you

Your voice grows with the database, not with ad spend. You approve what goes in; the stack publishes it to the profile and to the chat on your site; the re-audit measures how far the voice now reaches.

Tag count and questions answered are tracked in the dashboard, so growth of the voice is visible, not assumed.

Each requirement has an owner: a part of the stack and a person on the team.

The stack does the repeatable work. The team does the judgment. You see the result in one place.

REQUIREMENT	WHAT THE STACK DOES	WHAT THE TEAM DOES	WHERE YOU SEE IT
01 Readable facts	Site Intake captures every page and extracts baseline facts. The audit runs 22 technical checks and ships FAQ JSON-LD, llms.txt and an install pack for your own domain. AI Profile publishes one verified sentence per fact.	Verifies each fact with you before it goes live. Nothing unconfirmed is published.	Your AI Profile on aeofacts.ai and the chat on your site
02 One truth everywhere	The audit ranks every domain AI cites, shows where competitors out-cite you, and traces each error to its source. Site Intake flags every conflict with third-party listings.	Writes the fix briefs with the exact corrected value and the platform to change it on; sequences them in a 30/60/90 plan.	Gap list in the audit report
03 Answers to real questions	The AI Chat on your site logs every question it cannot answer. The audit asks 150 customer questions. The Learner asks you the questions customers put to AI, a few a week. GEM with AI captures what customers ask in person.	Builds the 150-question set with you, in your customers' words. Writes 12 articles a year, 3 per quarter, from the briefs and the unanswered questions; you approve and download them in the Learner (from November).	FAQ on the profile; 3 articles a quarter in the dashboard from November
04 Reasons to choose you	The audit benchmarks 3–4 named competitors on the same questions: territory by cluster, what AI associates with each, the questions they win. The profile carries an only-here section.	Writes the only-here features from the audit gaps and what customers value.	"Only here" on the profile
05 Kept current	The Learner keeps proposing; every approved tag publishes back to the profile and the chat automatically.	Ongoing tag review with you, so the profile moves with the business.	Tag history in the dashboard
06 A score, not an opinion	AEO Audit tool: 150 questions × 2 engines × 5 runs. Mentions, citations, share of voice vs competitors, by cluster and engine, with confidence intervals. Re-run unchanged.	Reads every mention, classifies errors into six types with root cause, and briefs you on what moved and why.	Baseline vs re-audit in the dashboard

Build the foundation once. Run the loop continuously. Prove it each cycle.

AEO is not a launch; engines change and so do the questions. The work is designed to keep going without becoming a project every time.

FOUNDATION · ONCE



- ✓ Site Intake: every page captured, baseline facts, third-party gaps
- ✓ Baseline audit: 1,500 answers scored, so there is a number to beat
- ✓ Question set built with you, in your customers' words
- ✓ AI Profile verified and published

COVERS

Readable facts · One truth · A score

OPERATING LOOP · CONTINUOUS



- ✓ The Learner proposes tags from reviews, gaps and customer questions, including what the chat could not answer
- ✓ You confirm or reject in the dashboard; the tag database grows every week
- ✓ A few questions a week with the team, not a workshop a quarter
- ✓ Each new tag answers more questions, on the profile and in the chat; visibility tracked as it moves

COVERS

Real questions · Reasons to choose you · Kept current

PROOF · EACH CYCLE



The same 150 questions, the same engines, the same scoring, run again. The delta is the result.



COVERS

A score, not an opinion

Proven means all six requirements are checked in the dashboard, the tag database is larger than at baseline, and the re-audit score is higher on the identical test.

Every month, see what is happening **on your own site.**

A traffic tab in the AI Learner. It reports monthly, between the quarterly visibility audits.

COMING IN NOVEMBER



Website traffic, month by month

Growth from GA4 in one view, including visits GA4 attributes to AI assistants such as ChatGPT, Gemini and Perplexity.



You choose how the data arrives

Share GA4 access and we upload the report for you, or download it yourself and send it to us each month.



AI Profile traffic. COMING LATER

Visits to your AI Profile and click-throughs from it to your site, next to your own traffic.

AI LEARNER · TRAFFIC

ILLUSTRATIVE

Website sessions (GA4)

Visits from AI assistants



Pairs with the quarterly re-audit: the audit shows what AI says about you, the traffic tab shows who arrives. AI Profile traffic is added once tracking is ready.

Baseline in month one. Fixes in month two. Proof in month three.

Timings start when we receive your intake sheet. Optional items launch when you want them.

MONTH 1

Set up and baseline

ONBOARDING · WEEKS 1–2

- ✓ Intake sheet sent to you
- ✓ Content sheet for GEM with AI, your customer-facing web app, sent to you
- ✓ AI Profile built from your website; link shared for your sitemap
- ✓ GA4 access or monthly report requested for the traffic dashboard

WITHIN TWO WEEKS OF YOUR INTAKE SHEET

- ✓ 150-question set built in your customers' words
- ✓ AEO audit and AI visibility report delivered
- ✓ First article questions sent to your team
- ✓ First 12 questions for new AI Profile facts
- ✓ Crawlability files for your website (llms.txt, FAQ schema)

MONTH 2

Fix and launch

- ✓ AI Profile fully reviewed by a team member
- ✓ Fixes recommended for the errors AI makes about you, with the pages to correct
- ✓ First 3 articles (once your article questions are answered)
- ✓ Next 12 questions; new facts added to your AI Profile
- ✓ AI Chat launched OPTIONAL
- ✓ GEM with AI launched (within two weeks of your content sheet) OPTIONAL
- ✓ Questions AI Chat could not answer added to your dashboard as they come in
- ✓ GEM with AI reviews tracked and added to the Learner
- ✓ Crawlability re-tested to confirm the fixes

MONTH 3

Verify and measure

- ✓ Third-party listing check: what booking sites and directories say about you vs your website, plus facts missing from your site
- ✓ Next 12 questions; new facts added to your AI Profile
- ✓ AI visibility audit re-run on the identical 150 questions
- ✓ First quarterly snapshot against your baseline
- ✓ AI Chat questions and GEM with AI reviews keep feeding the dashboard and Learner

ONGOING Fix recommendations follow one month after each quarterly visibility audit.

SECTION THREE

How we do it: the tech stack.

Collect, measure, fix, publish, learn. Each product, what it does and what you get from it.

01

AEO Audit tool

02

AI Profile

03

Site Intake tool

04

AI Learner

05


AI Chat

06

GEM with AI

Everything feeds the **Learner**. The Learner grows the profile. The profile gets you found.

Re-audit on the same test to measure the gain



AEO Audit tool

Asks ChatGPT and Gemini 150 customer questions, five times each: 1,500 recorded answers about a business and its competitors.



WHAT IT IS

An in-house pipeline on the OpenAI and Gemini APIs. In: a certified 150-question set, frozen so it re-runs identically. Out: a structured, auditable run.



WHAT IT DOES

150 questions × 2 engines × 5 runs. Scores mentions, citations and share of voice against named competitors, by cluster and engine; every cited domain; 22 technical checks.



WHY IT EXISTS

The baseline everything is measured against, and the first collector: every run exposes the gaps that become tags.

AEO Tool by Tabi Life · AI Visibility Dashboard 2025-10-21 14h 50p.00

Internal use only — sampled AI visibility estimate, not production data.
Engine coverage is partial (ChatGPT + Gemini only, via the mock engine in this demo). Scores are a sampled share-of-voice estimate from a fixed prompt set — not a literal count of real AI user queries. Not production-ready: live adapters, secrets, storage, auth, and compliance are still pending.

Run an instant audit

Website URL * **Business name**

What do they do? (helps generate better questions) — business type + area + key features + target guests, as one phrase

Competitor websites (comma-separated URLs, optional)

Question language
Auto-detect from website (jp → Japanese)

Market — where the simulated customer searches from (choose live engine results)

Report language — locale of the report + deliverables
Same as question language

☒ **Save Demo (no API keys — uses mock engine)**

Estimate cost **Run now** Auto-generates the question set, runs the engines, builds the report.

Compare vs competitors (benchmark)

Website URL **Your business name**

Run output as surfaced in the AI Learner (demo property, fictional data)

The audit report: a full diagnosis, with up to 15 briefs.



Scorecard and competitors

Mention, citation and share of voice against 3–4 named competitors, by cluster and engine.



Where AI reads about you

Every cited domain ranked; where competitors out-cite you; each error traced to its source.



Every mention reviewed

Analysts read all mentions; six error types; root cause: third party, engine or site.



Up to 15 briefs, 30/60/90 plan

Third-party fixes with corrected values, page rewrites, FAQ JSON-LD, install pack. Re-audit on the same test.

デモページです。この施設は存在しません。AI Profile の掲載にリンク先確認のための架空データを表示しています。

aeofacts.ai · AI PROFILE JA EN 公式サイト

日本・神奈川県・横浜市中区・馬車道・タビライフホテル 横浜支店
ビジネスホテル

タビライフホテル 横浜支店 Tabi Life Hotel Yokohama

たびらひほてるよこはま・Tabi Life Hotel Yokohama・ホテル馬車道イン (～2023年)
ビジネスホテル・87室・神奈川県横浜市中区馬車道4-XX

✓ 施設により確認済み・タビライフ AI Profile 監査 ● 更新中のプロフィール 85件の事実を確認済み・掲載記事 61-618

報道日: 2026-08-01 次回監査予定: 2027-02-01 施設による最終確認: 2026-08-12

タビライフホテル 横浜支店は横浜市中区の馬車道エリアにあります。タビライフホテル 横浜支店は馬車道駅（みなとみらい線）5番出口から230m、徒歩3分です。タビライフホテル 横浜支店の客室数は87室です。タビライフホテル 横浜支店には人工温泉の大浴場があります。タビライフホテル 横浜支店のチェックインは15:00です。

チェックイン: 15:00	チェックアウト: 11:00	住所: 神奈川県横浜市中区馬車道4-XX	馬車道駅（みなとみらい線）: 5番出口 230m・徒歩3分
客室数: 87室	大浴場: あり（人工温泉浴）	駐車場: 22台 車庫利用 2.1m	支払い方法: 現金・Visa・Mastercard・ JCB・Amex・UnionPay 交通系IC・PayPay・Alipay・WeChat Pay
対応言語: 日本語・英語 24時間 ペトアル編・中英語訳 90-2200	ペット: 不可（補助犬は可）	キャンセル制度: 前日15:00まで無料 後日10%・予定100%	電話: +81-45-000-0000
公式サイト: https://tabilife-	宿泊税: 1人1泊 200円（チェックイン ～）	Wi-Fi: 実測中央値 400Mbps（有線 ～）	喫煙: 全室禁煙（2階に喫煙ブース）

Audit fixes go live on your AI Profile (demo property)

Not a score. A diagnosis of how AI sees you, who it believes, and what to fix first.

Nine sections, 60 to 70 pages, every number traceable to a raw answer. Figures below are from a real baseline audit, anonymized.

25%

of 1,500 answers name the business

13%

cite the business's own site

#2 of 4

share of voice vs named competitors

41%

of mentions contain an error, each traced to its source

114/150

questions the site already answers

13

briefs this time, with JSON-LD, ready to implement



SECTION 1

Executive summary

Mention, citation, share of voice, accuracy, tone, technical score, site-answerability, and what it means for the business. Readable in two minutes.



SECTION 2

How AI sees you today

Question by question: which of the 150 you win, which you lose, the exact wording AI uses. Accuracy and tone: every mention read by an analyst.



SECTION 3

Competitor analysis

Three to four named competitors on the same test: share of voice, territory by cluster, what AI associates with each, the questions they win.



SECTION 4

Citation sources

Every domain AI cites, ranked. Where competitors out-cite you, and the platforms to fix first.



SECTION 5

Content gaps and briefs

Site-answerability for all 150 questions, then up to 15 typed briefs: optimize existing, new page, third-party fix, technical. Only what the findings justify. Each with priority, effort, target questions and FAQ JSON-LD.



SECTION 6

Technical and crawlability

22 checks: AI-crawler access, server-side text, JSON-LD, llms.txt, sitemap freshness. Fix pack included.



SECTION 7

30/60/90-day roadmap

Third-party fixes first, then existing pages, then new pages. Facts for the business to confirm. Baseline numbers, including the blank zone nobody wins, for the re-audit.



SECTION 8-9

Method and appendices

Frozen question set, counting rules, confidence intervals, every raw answer and every cited domain. Nothing in the report is unsourced.

Delivered as an A4 report in the client's language, with a five-page visual summary up front. Example figures: one business, one baseline audit, 2026.

AI Profile

Your verified facts on *aeofacts.ai*: one page that AI engines and people can read.



WHAT IT IS

A crawlable page for each business, in Japanese and English, built only from the facts you have approved.



WHAT IT DOES

One sentence per fact with its source, FAQ, only-here features and schema markup. 18 checks before anything goes live.



WHY IT EXISTS

AI engines need one clean, verified source per business. This page is that source, and it updates without a website release.

デモページです。この施設は存在しません。AI Profile の掲載にリンク先のための架空データを表示しています。

aeofacts.ai · AI PROFILE

JA EN 公式サイト

日本・神奈川県・横浜市中区・馬車道・タビライフホテル 横浜支店

ビジネスホテル

タビライフホテル 横浜支店 Tabi Life Hotel Yokohama

たぐりいふほてるよこはま・Tabi Life Hotel Yokohama・ホテル馬車道イン (～2023年)
ビジネスホテル・87室・神奈川県横浜市中区馬車道4-XX

✓施設により確認済み・タビライフ AI Profile 監査 ●更新中のプロフィール 85件の事実を確認済み・掲載基準 G1-G18

報道日: 2026-08-01 次回監査予定: 2027-02-01 施設による最終確認: 2026-08-12

タビライフホテル 横浜支店は横浜市中区の馬車道エリアにあります。タビライフホテル 横浜支店は馬車道駅（みなとみらい線）5番出口から230m、徒歩3分です。タビライフホテル 横浜支店の客室数は87室です。タビライフホテル 横浜支店には人工炭酸泉の大浴場があります。タビライフホテル 横浜支店のチェックインは15:00です。

チェックイン: 15:00	チェックアウト: 11:00	住所: 神奈川県横浜市中区馬車道4-XX	馬車道駅（みなとみらい線）: 5番出口 230m・徒歩3分
客室数: 87室	大浴場: あり（人工炭酸泉）	駐車場: 22台 車高制限 2.1m	支払い方法: 現金・Visa・Mastercard・ JCB・Amex・UnionPay 交通系IC・PayPay・Alipay・WeChat Pay
対応言語: 日本語・英語 24時間 ペタム線・中央線 6:00-22:00	ペット: 不可（補助犬は可）	キャンセル規定: 前日15:00まで無料 当日10%・平日100%	電話: +81-45-000-0000
公式サイト: https://tabilife-	宿泊税: 1人1泊 200円（チェックイン	WiFi: 実測中央値 400Mbps（有線	喫煙: 全室禁煙（2階に喫煙ブース）

aeofacts.ai profile page (demo property)

The page: what AI engines and customers read about you.



Verified header

Badge, audit date, next audit, last owner confirmation; key facts grid; link to the official site.



FAQ and only-here

Questions customers ask, answered from tags; the features only this business has, as chips.



One sentence per fact

Every tag as a plain statement with its field, value and an (i) showing source and tier.



Unconfirmed, stated

What is not yet confirmed and why, so AI does not guess. Contact block and data notice.

このページの内容 施設概要 設備 アクセス 客室 お風呂 食事 宿泊規定 よくある質問 この施設だけ 荷物・サービス 周辺 関連プロフィール

タビライフホテル 横浜支店の読み方は？
タビライフホテル 横浜支店のかな表記は「たびらいふほてる よこはま」です。①

タビライフホテル 横浜支店のローマ字表記は？
タビライフホテル 横浜支店のローマ字表記は Tabi Life Hotel Yokohama です（マクロンなし表記も同一施設）。②

タビライフホテル 横浜支店は何のブランドの施設ですか？
タビライフホテル 横浜支店はブランド「TLH」の横浜支店です。③

タビライフホテル 横浜支店の旧名称は？
タビライフホテル 横浜支店は2023年までホテル馬車道インという名称でした。④

タビライフホテル 横浜支店はいつ開業しましたか？
タビライフホテル 横浜支店は2016年に開業しました。⑤

タビライフホテル 横浜支店の耐震性は？
タビライフホテル 横浜支店は2019年に耐震補強工事を実施しました。⑥

▶ その他の質問

↑ タビライフホテル 横浜支店にしかない特徴は？

馬車道駅（みなとみらい線） → 5番出口 230m ・徒歩3分 馬車道駅（JR横須線） → 7番出口 350m ・徒歩4分

馬車道駅（横浜市営地下鉄ブルーライン） → 1番出口 400m ・徒歩5分 バリアフリー経路 → JR馬車道駅7番出口（エレベーターあり） 駐車場 → 22台

駐車料金 → 1泊 1,500円（要予約） EV充電器 → 6kW × 2基（500円） ユニバーサルトイレ 仕様 → 入口幅83cm・ロールインシャワー・引き出し押出

女性専用フロア → 7階 12室（カードキー制御） 大浴場の種類 → 人工温泉浴 大浴場（男性） → 15:00-24:00

大浴場（女性） → 6:00-9:30（入れ替え制） タトゥー → 利用可（カパー不要） 貸切風呂 → 45分 2,200円（22:00まで）

FAQ and only-here section (demo property)

Site Intake tool

Sitemap in, screenshots out, baseline tags next.



WHAT IT IS

A crawler: paste any site's sitemap and it opens every page and saves a full-page screenshot, named per page, to Drive.



WHAT IT DOES

Screenshots are read into a tag workbook: the official site's facts become the baseline tag set, with third-party listings aligned row by row to confirm or contradict them.



WHY IT EXISTS

The starting point of every AI Profile and the input to the gap analysis against third-party sources; conflicts go to the curation queue.

Works for any business with a website.

KEY CAPABILITIES



Sitemap → every page captured



Baseline tags from the official site



Third-party facts aligned row by row



Gaps and conflicts to the queue

The baseline: every page captured, every fact tagged.



Full-page capture set

One named screenshot per page of your site, filed to Drive: the evidence record for everything that follows.



Baseline tag workbook

Your site's facts extracted as searchable tags in English and Japanese: the first version of your profile.



Third-party comparison

The same facts aligned row by row against third-party listings: confirmed, missing or contradicted.



Gap and conflict list

What your site does not say, and where sources disagree, handed to the curation queue and the audit.

The screenshot shows the AI Learner interface for a property named 'タビライフホテル 横浜支店'. The left sidebar contains navigation options: AI Profile, AI Visibility, Curation Queue (highlighted), Questions, Guest Reviews, and Guest Demand. The main content area is titled 'All Tags — タビライフホテル 横浜支店' and includes a 'SAMPLE DATA' link. Below the title, there's a summary of tags and a table of statistics.

TAGS IN THE ACCEPTED SET — WHAT THE BUILD LAYER READS	OF THOSE, TAGS THE PUBLISH RULES ALLOW ON THE PAGE	PROPOSED TAGS AWAITING A DECISION IN THIS SERVICE	DISTINCT TAGS ON RECORD ACROSS BOTH
240	240	18	258

Below the table, there's a 'Profile' section with a 'PUBLISHED' status and a 'gate verdict for the current tags: not yet evaluated'. A note mentions 'The verdict is the build layer's own, shown as it reports it.' Below this, there's a section for '18 candidate tags awaiting confirmation' with a 'Go to the queue' link. At the bottom, there's a 'Filter by source type' section with options: Verified on-site, Confirmed across multiple sources, From public online sources, and Properly declared. The 'ACCESS' section shows a list of tags, including 'Airport hameda minutes' and 'about 40 min (Keiikyū → JR)'.

Baseline tags as they appear in the AI Learner after intake (demo property)

AI Learner Dashboard

Your control room: every fact, proposal and question in one place.



WHAT IT IS

A web app at portal.aefacts.ai where your facts live: review proposed facts, answer questions, see your AI visibility.




WHAT IT DOES

Takes in signals and your declarations; queues proposed tags for you to accept or reject; from November, 3 articles a quarter to download; publishes to the AI Profile.



WHY IT EXISTS

Keeps the profile alive and the tag set growing. Nothing reaches the public page without passing through this app.



All Tags view with source tiers (demo property)

The dashboard: what you see and do each week.



All tags, by tier

Every tag on record with its source: verified, corroborated, observed, declared. Filter, inspect, flag.



Curation queue

Proposed tags with the source excerpt. Accept publishes; reject keeps it off the profile.



Questions and reviews

2–3 questions a week driven by gaps and customer demand; what customers said in reviews via GEM with AI.



Visibility and articles

How AI answers about you today and, from November, 3 ready-to-publish articles a quarter to approve and download.



Curation queue (demo property)

AI Chat

Understands the question with ChatGPT, answers only from the AI Profile.



WHAT IT IS

A chat widget on your site. ChatGPT is used to understand the question; the answer itself comes only from your approved tag set and AI Profile.



WHAT IT DOES

1. Match the question to the profile. 2. Unclear? ChatGPT interprets the intent; the answer still comes from our database. 3. No fact? Say “please contact us” and log it to the Learner.



WHY IT EXISTS

To find where your voice is missing. Every unanswered question is a tag you do not have yet, asked by a real customer.

KEY CAPABILITIES



Understands with ChatGPT, answers from our data only



Shows the fact behind every answer



No fact? “Please contact us”, never a guess



Unanswered questions logged to the Learner

Works for any business with a website. ChatGPT interprets the question; it never writes the answer.

The chat: verified answers, and a map of what is missing.



Answers from the record

Customers ask in their own words; the answer comes from your AI Profile, with the fact shown.



Nothing invented

If no verified fact exists, the chat says “please contact us” instead of guessing, and logs the question.




Questions to the Learner

Every unanswered question queues in your dashboard as a tag candidate. You answer; we turn it into a tag.



Your voice, where it is needed

Unanswered questions show which tags to build next, so your voice grows exactly where customers asked.



AI Chat answering from the profile and logging an unanswered question (concept)

GEM with AI

Customer-facing web app · opens from a QR code, no download required.



WHAT IT IS

A QR and browser platform customers open on site. The 2026 renewal of our smartphone concierge (SPOT, 2021–2024).



WHAT IT DOES

Guide, mobile order, coupons, six-language chat. Collects reviews and sees what customers tap on: the raw material for what they will ask an AI.



WHY IT EXISTS

The on-site collector. Reviews and usage become observed tags and the questions the Learner puts to you.



Image: tabilife.co.jp

On site: what customers open, and what you learn.



On-site guide

Facility info, FAQs, services and local recommendations in the customer's language, branded to you.



Order, coupons, chat

Mobile order for in-house shops and dining, coupons and vouchers, six-language chat with staff handoff.




End-of-visit reviews

A short question at the end of the visit; scores by category against comparable businesses.



Usage insight

What customers tap on, by language and time, feeds observed tags and the questions they will ask an AI.



Customer reviews from GEM with AI, as shown in the AI Learner (demo property)

SECTION FOUR

Next steps.

Your questions answered, then where to start: a one-off audit, or a monthly package that keeps the loop running.

01

Common questions

02

AEO Audit · one-off

03

Core · monthly

04

Growth · monthly

The questions we hear most, answered.

What decision-makers ask before they start.

01

Why start now?

Few businesses work on this yet, so early movers set the facts AI learns. Facts you add keep working and build on each other; an ad stops when the spend stops.

02

Isn't structured data, consistent listings and more content enough?

Each covers one of six requirements. AI combines them, and if one is missing it hedges or names someone else. We run all six as one loop.

03

AI keeps changing. Will the work be wasted?

The details change; the basics don't: readable, consistent, current facts. Your AI Profile sits on our platform, so we update it as engines change.

04

What if our competitors score well too?

The score is the foundation. The difference is depth: only-here facts and answers to more customer questions. Every question you answer widens the gap.

05

How is this different from other AEO services?

Most sell one piece. We diagnose, deliver the fixes (llms.txt and FAQ schema included), publish a verified AI Profile, learn from real customer questions, and prove it on the same test.

06

What results can we expect?

AI names you when customers ask by need, not by name, reaching people who don't know you yet. Measured at every re-audit; expected ranges, not guarantees.

07

Aren't 150 questions too many?

150 covers every topic customers ask about, so the real priorities show. Frozen, so each re-audit is a fair comparison. New facts come separately: about 12 Learner questions a month.

08

How much work is it for our team?

An intake sheet to start, then a few questions a week in the AI Learner. We run the audit, build the profile and write the fixes; you approve facts and your web team installs the changes.

09

Does it help SEO too?

Much of the work overlaps: structured data, accurate and complete facts, clear FAQs. The goal differs: SEO ranks pages; AEO gets you named and cited.

10

Is it worth the investment?

It builds an asset, not ad spend: a verified fact base that keeps answering and grows each quarter. You judge the value with your own numbers.

Start with the audit. Choose how fast you move.

Pricing is set for your market and industry.



ONE-OFF

AEO Audit

Find out where you stand in AI answers today.

¥ _____ one-off

- ✓ 150 questions in your customers' words, asked across ChatGPT and Gemini, 5 runs each
- ✓ Mentions, citations and share of voice against the competitors you name
- ✓ Every mention reviewed; every error traced to its source
- ✓ Up to 15 briefs ready to implement, with the exact values to publish
- ✓ AI visibility report
- ✓ AI crawlability fix pack: llms.txt and FAQ schema, ready to install



MONTHLY

Core

Fix what AI gets wrong, then keep it right.

¥ _____ per month

EVERYTHING IN YOUR FIRST 90 DAYS, THEN:

- ✓ AI visibility re-audit every quarter
- ✓ Fix recommendations after each audit
- ✓ AI Profile and AI Learner dashboard
- ✓ 12 articles a year
- ✓ Monthly traffic dashboard (from November)
- ✓ Third-party listing check every year
- ✓ Crawlability files kept up to date



MONTHLY

Growth

Move faster, with every product included.

¥ _____ per month

EVERYTHING IN CORE, PLUS:

- ✓ AI visibility re-audit every month
- ✓ Fix recommendations every month
- ✓ 24 articles a year
- ✓ Third-party listing check every quarter
- ✓ AI Chat included
- ✓ GEM with AI included

[ADD TO CORE](#) AI Chat · GEM with AI

[MULTIPLE LOCATIONS](#) Volume pricing per location