

AN AI VISIBILITY REPORT • PREPARED BY OOMIRA

Independent assessment by Oomira, from public material. Not commissioned by, affiliated with, or endorsed by Stripe.

How AI sees Stripe

What assistants can reach about this company, what they say when asked, and what it would take to change either.

WHAT OOMIRA HOLDS ABOUT STRIPE

371

SOURCED FACTS

50

SOURCES CITED

35

SITES READ

48

PEOPLE

67

COMPANIES

102

DATED MOMENTS

REACHABILITY

Can a machine get to the facts you want a buyer to see?

78%

Mostly reachable

20 passing and 7 partial out of 30 checks.

HOW MUCH OF YOU A MACHINE CAN REACH

**You 78%**

Cloudflare 79%

SENTIMENT

Is what it says now more likely to win the deal or lose it?

Mixed**Read from the four assistant answers**

Praise and warnings in the same breath.

HOW THE ANSWERS READ

**You: Mixed**

Cloudflare: Mostly positive

Cloudflare is a best-in-class reference, not your industry. Industry benchmark: coming soon (run any competitor at oomira.com).

12 August 2026

Oomira • oomira.com

Scored 2026-08-12 • run Cb9BBNG_ • 30 checks counted

01 Executive summary

EXECUTIVE SUMMARY · GROUNDED IN THIS LIVE RUN

Assistants praise Stripe, then hedge on risk policies and freezes

THIS RUN 78% Mixed 20 of 30 checks working	IF YOU FIX IT The first three take this score to 85%. All nine: 98%. The projection is arithmetic over the same checks, not a forecast. What assistants say is not projected: they read the live web, so a shipped fix can land in the very next answer, on their clock.	WHAT THE SCORE MEASURES Ways an assistant can reach and use your public information today.
--	---	--

STATE OF THE WORLD

4 ANSWERS READ

All four assistants recommend Stripe unprompted, citing \$1.9T 2025 volume and \$159B February 2026 valuation. Two attach caveats about account freezes, pricing, support and automated risk decisions.

REACHABILITY

20/30 WORKING

stripe.com is fetchable as prose, admits every answer-time and training crawler, has a sitemap, Wikipedia and Wikidata entries. But the top ten results for a buyer query sit on surfaces Stripe doesn't own.

CONVERSION SIGNAL

2 OF 4 CAVEATED

Mixed. Buyers get a strong default recommendation, but Claude's freeze-and-support criticism and ChatGPT's "build redundancy" condition insert a due-diligence step before Stripe is chosen.

Criticisms are fair too: pricing is premium versus raw processors, account freezes frustrate small merchants, and support can feel thin.

CLAUDE · VERBATIM

WHAT STANDS OUT

- Answers cite third-party sources — Bloomberg, PYMNTS, Wikipedia, Crunchbase — over Stripe's own pages; zero of ten top results are owned surfaces.
- docs.stripe.com serves clean markdown and a copy-for-LLM button; stripe.com's marketing pages serve neither.
- The live MCP endpoint at mcp.stripe.com answers, yet nothing on stripe.com declares it exists.

WHAT TO DO NEXT

1. Publish an account-risk and freeze policy page with dated appeal timelines and support SLAs on stripe.com, not just docs.
2. Extend the docs markdown and copy-for-LLM treatment to stripe.com marketing and customer pages.
3. Declare the live mcp.stripe.com endpoint and the existing feed from stripe.com so agents can discover both.

Then rerun: Re-ask the same four assistants the buyer query used here — Stripe 2026 valuation and payment volume, plus a "should I use Stripe" prompt — after the risk-policy and support pages publish, and compare cited URLs.

02 What four assistants said

Two of the four attached a warning to the answer.

Mixed PRAISE AND WARNINGS IN THE SAME BREATH

Read from the four answers themselves. Two of the four raised one reservation each. Those cautions are what set it.

Each of these is the answer verbatim, exactly as the assistant gave it after its own search. Nothing is trimmed to make a point.

HOW TO READ THE MARKS

A passage marked like this is where an assistant warned a buyer off. The number ties it to the reading it was counted from, below. Nothing inside a mark is edited: the marks are drawn around the assistants' own words, whole.

CLAUDE WARNS THE BUYER

Stripe is one of the most impressive infrastructure companies of the past 15 years. Developer experience is its real moat, clean APIs, excellent docs, and the numbers back it up: businesses on Stripe processed \$1.9 trillion in 2025, up 34%, roughly 1.6% of global GDP, with a \$159 billion valuation in a February 2026 tender, up from \$106.7 billion.

Criticisms are fair too: pricing is premium versus raw processors, account freezes frustrate small merchants, and support can feel thin.¹ Its stablecoin and wallet bets (Bridge, Privy, Tempo) are interesting but unproven. Still no IPO, notably.

WHAT CLAUDE LOOKED AT

SEARCHED Stripe 2026 valuation payment volume
OPENED [Fintech investments on pymnts.com](#) · [2026 02 24 on bloomberg.com](#)

CHATGPT WARNS THE BUYER

Stripe is one of the strongest payments platforms for internet businesses, especially developers, SaaS companies, and marketplaces. Its APIs, documentation, checkout tools, billing, fraud controls, and Connect product substantially reduce the effort of accepting and routing money. It has also grown beyond card processing into broader financial infrastructure; Stripe says businesses on it processed \$1.9 trillion in 2025. (stripe.com) My reservation is concentration risk: payments are mission-critical, and automated compliance or risk decisions can create serious disruption.² Compare total cost, payout reliability, support, dispute handling, and account-risk policies against alternatives. It is an excellent default choice, provided you build redundancy.

WHAT CHATGPT LOOKED AT

SEARCHED [site:stripe.com](#) Stripe products payments billing connect treasury official
OPENED [Stripe 2025 update on stripe.com](#)

PERPLEXITY

NEUTRAL

Stripe looks like one of the strongest payments infrastructure companies: it provides developer-friendly tools for payments, billing, fraud prevention, and money movement, and it is used by businesses of many sizes.[2][7][8] Its scale is impressive too, with reports placing it at about \$1.9 trillion in payment volume in 2025 and a valuation around \$159 billion in 2026. [1][3][13] I'd describe it as a high-quality, enterprise-grade platform that is especially attractive if you want flexibility and global reach. The main tradeoff is complexity: Stripe is powerful, but some businesses may find its pricing and product depth more than they need.[15]

WHAT PERPLEXITY LOOKED AT

OPENED

Stripe, Inc on en.wikipedia.org · stripe.com · stripe on komo.ai · +14 more

GEMINI

SPEAKS WELL OF YOU

Stripe is an exceptional technology company that has fundamentally reshaped internet commerce. Known for its developer-first approach, it offers elegant, highly customizable APIs that make integrating complex payment processing remarkably simple for businesses of all sizes, from budding startups to global enterprises.

Beyond basic payments, Stripe's expansive ecosystem now includes powerful tools for billing, fraud prevention, corporate cards, and company incorporation. Its relentless focus on innovation, reliability, and user experience has made it the gold standard in financial infrastructure. Ultimately, Stripe succeeds by eliminating the friction of digital transactions, serving as a vital engine for the global internet economy.

SEARCH RECEIPT

This assistant returned no retrieval receipt.

Where a buyer hesitates

Assistants flag account freezes, premium pricing, thin support and automated risk decisions, urging buyers to compare alternatives and build payment redundancy first. Two of the four raised one reservation each. Each is highlighted in the answers above, in their own words.

CONTENT THAT RETIRES THE OBJECTIONS

Each recommendation starts with a buyer hesitation above, then names the public evidence that should replace it.

BUYER
HESITATION

Automated compliance or risk decisions can freeze accounts and cause serious disruption (ChatGPT, Claude).

PUBLISH

A dated account-risk and reserves policy page: what triggers a hold, appeal path, resolution timelines, payout continuity commitments.

PLACE IT

stripe.com top-level page linked from /pricing and the sitemap, with FAQ markup like /pricing already carries.
Assistants already fetch and quote stripe.com prose; a dated, question-shaped page gives them a citable answer instead of a caveat.

BUYER HESITATION PUBLISH PLACE IT	Support can feel thin, especially for small merchants (Claude). A support tiers page stating response targets, channels and escalation by plan, with the date printed. stripe.com/pricing, alongside the published plan figures, in the same FAQ markup already present there. <i>That page is the one plan surface agents read in figures; a support answer beside the price is lifted verbatim.</i>
BUYER HESITATION PUBLISH PLACE IT	Pricing is premium versus raw processors and product depth may exceed what a business needs (Claude, Perplexity). A total-cost comparison page: effective rate by volume band and which products a small merchant can skip. stripe.com/pricing with Offer and PriceSpecification markup, currently absent from that page. <i>The price is readable only as design; structured offers plus a cost breakdown let an agent answer the cost objection numerically.</i>

The readings behind Mixed

Three readings that back it up, each one counted from something on this page rather than assigned. The word above is read from the answers themselves.

■ PASS Confidence 2 observations 2 assistants gave a buyer a material caution before committing. That is part of the answer they receive, not an Oomira warning. COUNTED FROM The sentences marked 1 and 2 in the answers above, from Claude and ChatGPT. Every word of it is theirs.	■ PASS Proof 0 issues found The answers carry checkable evidence: named work, and somebody other than you attesting to it.	■ PASS Accuracy 0 issues found Nothing said about you is wrong or out of date. COUNTED FROM Every claim above that your record can test, and not one of them differs from it.
---	--	---

All three readings behind it come back clean.

03 What the 30 checks found

THE AUDIT MAP

Every square is one check we ran today. The audit below writes every check out: what we checked, what came back, the evidence, and the next change where one is needed.

20/30
CHECKS
WORKING

WORKING PARTIAL · COUNTS HALF NOT WORKING WARNING · UNSCORED

ROUTE	CHECKS · ONE SQUARE EACH	SCORE
1 Search about you top results, asked like a buyer		1/1
2 Your site to machines what a program receives from your pages		8/9
3 The machine front door how agents discover and call you		2/6
4 What an agent can buy price, plans and actions a machine can use		3.5/5
5 Who else carries you press and independent proof		4/4
6 Knowledge graph Wikidata, Wikipedia and entity records		2/2
7 Social and content surfaces profiles and content surfaces machines can read		3/3

THE FULL AUDIT		EVERY CHECK, WRITTEN OUT
CHECK	RESULT & EVIDENCE	NEXT CHANGE

ROUTE 1			1/1
Search about you			
Searched the way somebody asking about you would WORKING Whether a program that visits your site gets your actual words back. Your browser runs JavaScript to build the page; most machines do not. If the words only appear after that code runs, a machine sees an empty page where you see a homepage.	10 of 10 top results are about Stripe (query: “Stripe 2026 valuation payment volume”), 0 of them on a surface you own. 1. about you · Stripe statistics on chargeflow.io 2. about you · 2026 02 24 on bloomberg.com 3. about you · Stripe statistics on capitaloneshopping.com 4. about you · Stripe statistics on axis-intelligence.com 5. about you · stripe on sacra.com 6. about you · blog on capitaly.vc 7. about you · Fintech investments on pymnts.com 8. about you · Stripe users on backlinko.com 9. about you · beststartup.us 10. about you · techcrunch.com	HOLDING · NO CHANGE NEEDED	
ROUTE 2			8/9
Your site to machines			
Fetchable as prose WORKING Whether a program that visits your site gets your actual words back. Your browser runs JavaScript to build the page; most machines do not. If the words only appear after that code runs, a machine sees an empty page where you see a homepage.	An anonymous fetch of stripe.com, no browser and no JavaScript, comes back as real prose.	HOLDING · NO CHANGE NEEDED	
Inner pages readable WORKING The same question, asked of the pages that carry your claims rather than the homepage. Marketing homepages are usually built to be readable; blogs, case studies and docs often are not.	3 of 3 inner pages sampled from stripe.com return real prose to an anonymous fetch, not just the homepage. /customers/dena /customers/venue /customers/wayflyer	HOLDING · NO CHANGE NEEDED	
What an AI agent actually receives WORKING What your server does when the specific programs behind ChatGPT and Claude come to read a page. They identify themselves by name, and a lot of security tooling turns them away without anyone deciding to. You cannot see this from a browser, because your browser is not one of them.	stripe.com was fetched as ChatGPT-User and served normally (HTTP 200), the same as a browser got at the same moment.	HOLDING · NO CHANGE NEEDED	

robots.txt to answer-time agents ■ WORKING robots.txt is a plain text file at the root of your site that tells automated visitors what they may read. Answer-time agents are the ones that fetch a page at the moment someone asks an assistant about you, as opposed to the ones that collect text to train future models. This asks whether that file lets the first group in.	stripe.com/robots.txt admits all 6 answer-time agents checked, the ones that fetch pages when someone asks a model about Stripe today.	HOLDING · NO CHANGE NEEDED
robots.txt to training crawlers ■ WORKING	stripe.com/robots.txt blocks none of the 5 training crawlers checked, so stripe.com's pages can enter future model weights.	HOLDING · NO CHANGE NEEDED
How much of the site agents may read □ WARNING · UNSCORED The same file, read for what it closes rather than whether it opens. It can admit a visitor and then place whole sections off-limits, which is easy to do by accident and impossible to notice.	robots.txt closes 8 paths to the answer-time agents, /bitcoin/refund, /handoff, /handoff-healthcheck, /sources/refund, +4 more, and not one of them carries a claim about Stripe: they are endpoints, fixtures and health probes. Nothing to fix, and not counted against you. The whole rule is below so you can check that reading. <pre> /bitcoin/refund /handoff /handoff-healthcheck /sources/refund /sources/sepa_mandate /sources/test_klarna /sources/test_source /unsupported-browser </pre>	NO EXACT CHANGE PRESCRIBED
schema.org identity markup ■ PARTIAL · COUNTS HALF A small block of invisible code in your page that spells out, in a format every machine already agrees on, "this page is about a company, its name is X, its website is Y, its founders are Z". Your page says all that in words a person can read; this says it in a way a machine cannot misread. It is roughly fifteen lines and it is what stops a search engine or a model attaching your page to the wrong company.	schema.org markup on the stripe.com homepage, and none on 3 of the 3 inner pages sampled: /customers/dena, /customers/venue, /customers/wayflyer. Homepage Organization markup ships with most templates; it is the markup on the pages carrying claims that decides which entity a ranker thinks a page is about.	Extend the markup to the pages that carry claims, Article or Product with an explicit entity reference, not just Organization on the homepage.
Dated claims ■ WORKING Whether the numbers on your site say when they were true. "We serve 40 clients" with no date is a sentence a model will either repeat forever or hedge around; "40 clients as of March 2026" is something it can quote.	3 claims on stripe.com carry a date the site itself printed. privacy policy last updated: 2026-04-28 traction summary: Businesses on Stripe generated \$1.9T in 2025. Stripe processed more than \$40B during Black Friday through Cyber Monday... 2025 01 01: Stripe, Processed more than \$40B for businesses during Black Friday through Cyber Monday 2025 while maintaining 99.999...	HOLDING · NO CHANGE NEEDED

Pages that answer a buyer question ■ PARTIAL · COUNTS HALF A page per question a buyer actually asks, where the page is the answer itself: the actual number, the actual argument, dated, in prose. This is not the SEO version of an answers page. A search engine ranks pages built to target a question; an assistant retrieves the sentences that answer it and quotes them verbatim. Write it the way you would answer the buyer in an email, because those sentences are what get lifted.	1 page on stripe.com carries buyer questions in FAQ markup, /pricing, so an assistant can lift the question and your short answer verbatim. Half credit, not full: an FAQ entry is two sentences an agent can quote, while a dedicated answer page, one question, the full argument, its own URL, is what gets retrieved and cited AS the answer when a buyer asks it. Today whoever wrote that page answers for you.	Add one page per question a buyer actually asks: you versus a named alternative, the market you are in, the objection you always get. A page that answers a question in full sentences is what a model quotes verbatim.
Sitemap ■ WORKING A list of every page on your site, at a fixed address, so a machine can find pages you have not linked from the homepage. Most site builders generate it automatically.	At least 498 stripe.com pages are listed, read from the first slices of a deeper index, so an agent can discover what else Stripe publishes.	HOLDING · NO CHANGE NEEDED
ROUTE 3 The machine front door 2/6		
A declared way in for agents □ NOT WORKING A small file at a fixed address that tells a program arriving at your domain who you are and what it can do with you, rather than making it guess from your marketing copy. It is a static file you write once.	This one is for every company, including yours, it is an identity file, not a developer feature. Nothing at stripe.com/.well-known/agent-card.json or /.well-known/ai-plugin.json today. They are the first addresses an agent tries, and they hold what it most needs: a plain statement of who Stripe is, what you do, and where everything else lives. Without one, every assistant describing you is working from your homepage copy.	Publish an agent card at /.well-known/agent-card.json: who you are, what you offer, and the addresses of anything an agent can call. It is a static JSON file, it costs an afternoon, and it is the difference between an assistant describing you and an assistant transacting with you.
A catalogue of what you expose □ NOT WORKING	You have developer documentation, so there is something of Stripe's an agent could call, but no catalogue at stripe.com/.well-known/api-catalog, no OpenAPI description at /openapi.json, and no agent-useful relations in your Link headers. An agent asked to integrate Stripe has to guess at URLs from prose, and mostly gives up.	Publish an API catalogue at /.well-known/api-catalog listing every API you expose and where each is documented. This is the one file that turns "they have docs somewhere" into "an agent can find and call this".
A live tool endpoint an agent can call ■ PARTIAL · COUNTS HALF	Your MCP endpoint at mcp.stripe.com is live and answering, and nothing on stripe.com declares that it exists, there is no server card at /.well-known/mcp/server-card.json. An assistant can only connect if a person already knew the URL and typed it in. The hard half is built; the file that makes it findable is not.	Publish a server card at /.well-known/mcp/server-card.json. The server already works, this is the file that lets an assistant find it without a human pasting the URL, and it is the cheapest distribution you have.

A plain-text version of every page ■ PARTIAL · COUNTS HALF Whether your pages can be served as plain text instead of a designed web page. When someone pastes your page into an assistant, they paste the whole thing: navigation, cookie banner, footer. A plain-text version means the argument arrives instead of the furniture.	docs.stripe.com serves a clean markdown representation of its pages (Accept: text/markdown) and stripe.com itself serves none. Everything under the documentation is already in the shape an assistant wants; the pages carrying your argument to a buyer are not, and those are the ones people paste.	Extend the markdown representation you already serve to the rest of the site. The mechanism works; it is only wired to some routes.
A copy-for-AI button on the page ■ PARTIAL · COUNTS HALF The human half of the row above: a button that copies the page in that clean plain-text format, so a reader can give the page to their assistant without thinking about it. The format is what the machine consumes; the button is how a person finds out it exists.	Copy for LLM, View as Markdown on docs.stripe.com, and nothing of the kind on stripe.com itself. A reader in your documentation is told they can hand the page to an assistant; a buyer on your marketing pages is not.	NO EXACT CHANGE PRESCRIBED
A feed a machine can follow ■ PARTIAL · COUNTS HALF RSS: a plain list of what you have published, in date order, that anything can subscribe to. It is how a machine learns you posted something without checking your site every day.	Stripe's feed is live, and nothing points at it: no <link rel="alternate"> in the page head and no mention in robots.txt, so a crawler following the standard finds nothing. The feed works if you already know the address.	The feed exists, declare it. One <link rel="alternate" type="application/rss+xml"> in the head of every page, and a Sitemap-style reference in robots.txt. This is two lines, not a project.
llms.txt □ WARNING · UNSCORED A proposed file that would tell AI systems how to read your site. Included for completeness and deliberately not scored: no major AI company reads it today.	Live at stripe.com/llms.txt, though there is no llms-full.txt beside it, the index exists, the content it indexes does not. Worth knowing: no major AI provider reads this file today, so this is a bet on where things are going rather than a channel that is working now. It is not counted for or against you.	NO EXACT CHANGE PRESCRIBED
Content Signals □ WARNING · UNSCORED	stripe.com/robots.txt carries no Content-Signal directives. That is a position, not a gap, the file declares how you are willing to have your content used, and having no preference on record is a legitimate answer. Not counted against you.	NO EXACT CHANGE PRESCRIBED
Standards we did not score □ WATCHLIST · UNSCORED	Other scanners will mark you down for DNS-AID records, WebMCP, Web Bot Auth and x402. Two of those cannot be judged from outside, WebMCP needs your page's JavaScript to run, and Web Bot Auth is about what your server does with a signed request nobody has sent it. The rest are drafts with almost no adoption, and grading a founder on our forecast of which draft wins is not a number worth having. They are listed so you know they were considered, not skipped.	NO EXACT CHANGE PRESCRIBED

ROUTE 4			3.5/5
What an agent can buy			
A price an agent can read WORKING Whether an assistant can find out what you cost. If your pricing page says "contact sales", an assistant comparing three vendors reports numbers for the other two and a shrug for you.	Published in figures at stripe.com/pricing, so an assistant can quote what you cost.	HOLDING · NO CHANGE NEEDED	
A plan list an agent can read WORKING	Your plans are set out in readable figures at stripe.com/pricing.	HOLDING · NO CHANGE NEEDED	
Structured price and availability NOT WORKING The same fifteen-lines-of-invisible-code idea as schema.org above, applied to what you sell: price, currency, what is included. It is the difference between a machine reading your price and a machine guessing it from your layout.	The stripe.com page that sells plans carries no Offer or PriceSpecification markup. The price is in the design, not in the data, so a machine has to infer it from your layout.	Add Offer/PriceSpecification JSON-LD to the pages that sell: price, currency, availability, and what is included.	
A checkout an agent can complete PARTIAL · COUNTS HALF	No agentic-checkout protocol declared on stripe.com (/well-known/ucp, /well-known/acp.json). Your price is readable, so an assistant can recommend Stripe and hand the buyer over, it just cannot finish the purchase itself. That is a ceiling on what it does for you, not a wall.	These are young standards and adopting one is a bet, not a fix, but the half of it that always pays is the price and the catalogue above, which is what every assistant reads whether or not it can transact. Do those first, and revisit this when a protocol wins.	
Documentation an agent can be pointed at WORKING	Developer documentation at stripe.com/docs, so an agent asked to integrate you has something to read.	HOLDING · NO CHANGE NEEDED	

ROUTE 5

Who else carries you



4/4

Press

■ WORKING

Coverage of you by somebody else, on a site a machine can read. It counts for more than anything you publish about yourself, because it is not you saying it.

20 pieces about Stripe across 17 outlets: forbes.com, wired.com, entrepreneur.com +14 more

forbes.com, 27 facts
atouchofbusiness.com, 27 facts
redcarpetlife.in, 22 facts
gropedia.com, 21 facts
businessmenstory.com, 19 facts
thecompanycheck.com, 18 facts
settervc.com, 18 facts
tradersunion.com, 15 facts
wired.com, 12 facts
future.com, 8 facts
+ 7 more, counted in the result

HOLDING · NO CHANGE
NEEDED

Press releases

■ WORKING

pwc.com carried a release for Stripe.

HOLDING · NO CHANGE
NEEDED

Owned media

■ WORKING

A place you publish that you control and a machine can read: a blog on your own domain, a newsletter with a public archive. LinkedIn does not count, because nothing can read it.

stripe.com

HOLDING · NO CHANGE
NEEDED

Community mentions

■ WORKING

People discussing you in public places machines read heavily, like Reddit or Hacker News. Earned slowly and never planted.

1,017 Hacker News submissions link stripe.com.

HOLDING · NO CHANGE
NEEDED

ROUTE 6

Knowledge graph



2/2

Wikidata entity

■ WORKING

Wikidata is a free, public database of things and facts about them, and it is the table models look at to decide which company of your name you are. An entry is a short structured record: name, website, founding date, who runs it. You can create one yourself in an afternoon and it is the single strongest defence against another company with a similar name collecting the answers meant for you. It is NOT Wikipedia, which is an article with a fame test attached.

An entity matches Stripe. Queried directly against the Wikidata API, the machine-readable registry models resolve names through, and your defence against a namesake collecting your answers.

HOLDING · NO CHANGE
NEEDED

Wikipedia article ■ WORKING The encyclopedia article, which is a different thing from Wikidata above. It has a notability bar you have to earn through press coverage first, and one created too early gets deleted, publicly. Not having one is normal and is not really a failing.	An article matches Stripe. Queried directly against the Wikipedia search API, and the run's own sweep cites it (en.wikipedia.org).	HOLDING · NO CHANGE NEEDED
ROUTE 7 Social and content surfaces ■■■ 3/3		
LinkedIn □ WARNING · UNSCORED	linkedin.com/company/stripe · linked from stripe.com, so it is Stripe's · unreadable to machines: even held, it is a wall to machines, only search snippets leak out, so an AI cannot read what you publish there. Not counted against you: the wall is LinkedIn's, not yours. FOUND	SUGGESTION Mirror anything that matters there onto stripe.com or a surface an agent can read.
X □ WARNING · UNSCORED	twitter.com/stripe · linked from stripe.com, so it is Stripe's · partly readable: readable by humans and paid APIs only; a raw fetch gets a wall, search snippets quote it. Research checks it via the X API. Not counted against you: what a machine gets from X is X's limit, not yours. FOUND	SUGGESTION Post the claims that count with their numbers and dates, and mirror them onto stripe.com, the API and search snippets are the only machine routes in.
YouTube □ WARNING · UNSCORED	youtube.com/@stripe · linked from stripe.com, so it is Stripe's · partly readable: the page resolves but the videos are not machine-readable, titles and descriptions are all an AI gets. Research checks it via the YouTube API. Not counted against you: what a machine gets from YouTube is YouTube's limit, not yours. FOUND	SUGGESTION Put the claims in the descriptions and publish transcripts; the video itself is invisible to an AI.
GitHub ■ WORKING	github.com/stripe · linked from stripe.com, so it is Stripe's · machine-readable: fully machine-readable when public, readmes and code are fetchable, citable surfaces. FOUND	HOLDING · NO CHANGE NEEDED
Instagram □ WARNING · UNSCORED	instagram.com/stripehq · linked from stripe.com, so it is Stripe's · unreadable to machines: a wall to machines: no API for an agent, nothing readable past the login. Not counted against you: the wall is Instagram's, not yours. FOUND	SUGGESTION Mirror anything that matters there onto stripe.com or a surface an agent can read.

Facebook ⏏ WARNING · UNSCORED	facebook.com/StripeHQ · linked from stripe.com, so it is Stripe's · unreadable to machines: a wall to machines: no API for an agent, nothing readable past the login. Not counted against you: the wall is Facebook's, not yours. FOUND	SUGGESTION Mirror anything that matters there onto stripe.com or a surface an agent can read.
TikTok ⏏ WARNING · UNSCORED	Nothing for Stripe on your record, nothing linked from stripe.com, and nothing in this run's searches. Not counted for or against you: it is a wall to machines and the content is video besides, so an account you do not have costs you no reachability. NOT FOUND	NO EXACT CHANGE PRESCRIBED
Your blog ■ WORKING	stripe.com · on your record · machine-readable: machine-readable, the cheapest citable surface you own; a dated post is a citation you can create today. FOUND	HOLDING · NO CHANGE NEEDED
Your podcast ⏏ WARNING · UNSCORED	Nothing for Stripe on your record, nothing linked from stripe.com, and nothing in this run's searches. Not counted for or against you: audio is not machine-readable, so a show you never recorded is not reach you are losing. NOT FOUND	NO EXACT CHANGE PRESCRIBED
Mirrored where a machine can read it ■ WORKING	LinkedIn, X, YouTube, Instagram, Facebook are walled to machines, and Stripe also publishes somewhere one can read: GitHub, Your blog, writing on stripe.com. That is the surface an AI quotes Stripe from.	HOLDING · NO CHANGE NEEDED

PROMPTS, READY TO PASTE

A PDF cannot hand a block of text to an assistant, so the same report is published as plain text with every prompt this run produced in a fenced block, ready to paste: [this report as plain text on oomira.com](#).

AFTER THIS REPORT

This report is the first use of your record. Connect it to your AI and the next deck, diligence answer, update, or research task starts from these same sourced, dated facts.

Connect your AI → [oomira.com/mcp](#)

ABOUT THIS DOCUMENT

Oomira produced this report on its own initiative, as a demonstration of its AI visibility audit. It is based solely on publicly available material and automated observations of stripe.com recorded on the dates stated in the document. It was not commissioned by, prepared in collaboration with, reviewed by, or endorsed by Stripe, Inc. The Stripe name is used only to identify the subject of the assessment, and is a trademark of Stripe, Inc.; the Stripe logo has been deliberately omitted. Scores and verdicts reflect Oomira's own methodology as applied on the stated dates, and may not reflect the current state of stripe.com.