

SIRVEIL, INC.

Sirveil Exposure API

Legal Set — Version 2.0

Effective **25 August 2026**.

This document is the end-user license agreement offered on the AWS Marketplace listing for the Sirveil Exposure Verification API. It contains three instruments, which together form the agreement in the order of precedence set out in Section 3 of the Terms of Use:

1. API Terms of Use — version 2.0
2. Acceptable Use Policy — version 2.0
3. API Privacy Notice — version 2.0

The canonical live versions are published at sirveil.ai/legal/api-terms, [/api-aup](https://sirveil.ai/legal/api-aup) and [/api-privacy](https://sirveil.ai/legal/api-privacy). This PDF is a frozen copy of those documents as at the effective date above.

API TERMS OF USE — SIRVEIL EXPOSURE API

*Version 2.0 · Effective 25 August 2026 · supersedes the versions of 2026-08-21 and 2026-08-19 and the page previously published at sirveil.ai/scan-api/terms. **Basis:** the technical deep dive of 2026-08-24 (code at [landing/v34 @ 5ca896e](#), live OpenAPI 1.4.0, production data plane queried 24 Aug 2026), the Input Requirements Brief of 19 Aug 2026, and the Status of Record of 20 Aug 2026. Where this document and the code disagree, the code is the fact and this document is the defect. Reviewed and settled 25 August 2026. Every question carried by the previous version is decided and written out; nothing in this document is provisional.*

About the plain-English notes. Under most sections there is a short note explaining what that section means. The notes exist to help; they are not legal advice, they do not replace the formal text, and where a note and the legal text differ, **the legal text controls**.

Before the fine print

This is a deliberately complete commercial agreement, and it is written to protect Sirveil. It should be. We operate against well-funded adversaries in an industry that litigates.

But understand what it is *for*. The fence in Section 7 is not there to catch you out — it is there because the entire premise of this company is that personal information should be harder to weaponize, and an API that answers questions about people can be turned into a weapon by the wrong buyer. Everything else in here — the license in Section 12, the billing rules in Section 10, the honest indeterminate in Section 9 — is written to give you more than the industry standard, not less. We would rather tell you exactly when we charge you and exactly what we cannot promise than win an argument later on wording you never read.

You will also find, in Sections 6.1, 9.1, 11.1 and 15.1, a plain list of things we do not have, do not check, and will not promise: the certifications we do not hold, the vendors that necessarily see an identifier in clear text, the numbers in our own responses that you should not build a threshold on, and the limits that can refuse your call. None of that is in here because a lawyer made us. It is in here because you would find it in diligence anyway, and a buyer who learns it in month six is a buyer we mis-sold.

1. Parties, and what this covers

These API Terms of Use ("**Terms**") are between **Sirveil, Inc.**, a Delaware corporation ("**Sirveil**", "we", "us"), and the business or professional entity that opens an API account or purchases access through a marketplace ("**Customer**", "you").

They cover the **Sirveil Exposure API** (the "**API**" or "**Service**") and nothing else. The Sirveil consumer application has its own separate Terms of Service and Privacy Policy, which do not govern API use and are not modified by these Terms.

2. Definitions

- **Verification** — a call to `POST /api/v1/verify`, and the billed SKU of that name.
- **Scan** — a call to `POST /api/v1/scan`, and the billed SKU of that name. *(These are the names that appear on your invoice. Sales conversations may say "check" and "sweep"; the contract and the invoice line say verification and scan.)*
- **Answer** — the response the Service returns to a billable call: for a Verification, the three-state verdict together with its evidence and explanation; for a Scan, the ranked, tiered, cited exposure cards together with the coverage block.
- **Billable Answer** — an Answer that Sirveil delivers, or makes available for retrieval, in response to a properly authenticated and well-formed call. An Answer is a Billable Answer whether it was computed fresh or served from

cache, whether the verdict is favorable to you or not, and whether or not you retrieve it. **Not a Billable Answer:** a call refused before any lookup occurs; a call rejected as malformed or unauthenticated; and an Answer Sirveil could not serve because of a fault on Sirveil's side or in an upstream provider (an "**Unservd Call**"). (See Section 10 for how each case bills.)

- **Documentation** — the API documentation Sirveil publishes for the Service, including the OpenAPI document the Service serves, as updated.
- **Identifier** — a data element about a natural person that you submit in a call.
- **Output** — the Answers, evidence, reports, and other material the API returns to you.
- **Rate Card** — Sirveil's published per-call prices for the Service, as in effect from time to time: the prices published in the Documentation and, for a purchase made through a Marketplace, the prices published on that Marketplace listing. Prices are not reproduced in these Terms, so that a price change is not an amendment to them.
- **Registry** — Sirveil's maintained list of domains and sources, against which findings are **classified** and over which coverage is **reported**. The Registry is not the list of domains a Scan queries. See Section 6.
- **Subject** — the natural person an Identifier relates to.
- **Marketplace** — AWS Marketplace, or another reseller or marketplace through which Sirveil makes the Service available and through which you purchase access. Section 3.1 states which channels are actually available.

3. The agreement, and which document wins

The agreement between us consists of, in descending order of precedence:

1. a signed written enterprise agreement between you and Sirveil, if one exists;
2. these Terms;
3. the **Acceptable Use Policy** (AUP), which is incorporated by reference and is part of these Terms;
4. the **API Privacy Notice**, incorporated by reference;
5. the Documentation and the Rate Card.

Marketplace purchases. If you buy access through a Marketplace, that Marketplace's own terms govern the *transaction* — ordering, invoicing, metering submission, taxes, refunds, and entitlement — and control over these Terms to the extent of a direct conflict on those subjects only. These Terms govern *use* of the Service in all other respects. Where a Marketplace requires Sirveil to offer an end-user license agreement on the listing, these Terms are the document Sirveil offers there, and the version displayed on the listing when you subscribe is the version that governs your purchase. Any refund or cancellation policy published on that listing applies to the transaction and is in addition to anything in these Terms; nothing in these Terms narrows it.

You accept these Terms by accepting them at account creation or key issue, by purchasing through a Marketplace, or by calling the API — whichever happens first.

***In plain English.** If we have signed something with you, that wins. Otherwise this document wins, and the AUP and privacy notice are part of it. If you bought through a marketplace, their rules cover the money — ordering, invoicing, refunds — and ours cover the usage.*

3.1 Channels through which the Service is sold

The Service is available for purchase through **AWS Marketplace**. That is the whole list today.

Sirveil's published specification, documentation, and integration artifacts may name other channels — a marketplace whose listing is built but not selling, an API aggregator, or a direct checkout — because they exist in code. **A channel named in an artifact is not an offer, and is not available until it is listed in this Section.** Do not build a purchasing integration against a channel that is not listed here; if you have, tell us at support@sirveil.ai before you rely on it.

Sirveil may add or remove a channel. Removing a channel is a change to how the Service is sold, not a reduction of the Service you have already bought; if it affects an account you are actively paying for, Section 11 applies.

***In plain English.** You can buy this on AWS Marketplace. If you see another store's name in our spec, it is scaffolding, not a shop. We would rather say that here than let you find out at the checkout.*

4. Who may use the Service

The Service is offered only to businesses and professional users, only in the United States, and only to persons who can form a binding contract. It is **not** offered to consumers for personal use, and it is not offered to anyone barred from receiving US goods or services under sanctions or export-control law.

You represent that you are not located in an embargoed jurisdiction, are not on a US government restricted-party list, and will not make the Service available to anyone who is.

***In plain English.** This is a business tool. Consumers who want their own exposure checked should use the Sirveil app, which is built for exactly that and priced for it.*

5. Accounts and keys

- Access requires an account and a secret API key. **Keys are bearer credentials:** anyone holding one can spend your money, so treat one like a password.
- A key is displayed **once**, at issue. Sirveil stores only a cryptographic hash of it, together with a short non-usable prefix for identification, and cannot recover, re-display, or email you an existing key. If you lose one, issue a new key and revoke the old one.
- An account may hold up to **ten unrevoked keys**. Revoking a key does not delete its history: usage already recorded stays attributed to the key that produced it, which is what makes an invoice auditable.
- All use under your key is your use, and you are responsible for it, including use by your personnel and contractors and by end users reaching the Service **through your application**. That does not permit sharing a key with another legal entity.
- One key, one organization. Keys may not be shared across legal entities, sold, sublicensed, or embedded in a client application where an end user can extract them.
- Tell us immediately at support@sirveil.ai if a key is exposed. We will revoke it, and you can revoke and reissue keys yourself from the console at any time. We are not responsible for charges incurred before you tell us, except to the extent the law says otherwise or the exposure was our fault.
- **Use-case attestation.** At key issue, and on request afterwards, you will tell us truthfully what you are using the Service for. We **may** require any account to certify its use case in writing, and may condition, suspend, or refuse access on that certification or on a refusal to give one. We do not undertake to require certification from every account, and not requiring it from you is not a waiver of the right to require it later, nor a representation that we have vetted any other account.

***In plain English.** The key is the password and we cannot get it back for you — that is deliberate, because a key we could recover is a key someone could talk us out of. Ten keys per account, revoke them yourself, tell us fast if one leaks. And expect us to ask what you are building; the honest answer is always the cheap one.*

6. What the Service does — and what it is not

The Service answers one question — *does this identity appear in the sources we can reach, and where* — in two shapes.

A Verification (POST /api/v1/verify) asks whether an identity you name is indexed on **one domain you name**. It is one search-index question, answered in three states, with the evidence and an ordered plain-language explanation of what was queried and what matched.

A Scan (POST /api/v1/scan) sweeps for one identity across multiple source lanes and returns ranked, tiered, cited exposure cards together with a **coverage block** that names what we did not reach.

What a call must carry. Both endpoints require `firstName` and `lastName`; a Verification also requires a `domain`. That is not form validation — the name is the anchor every candidate record is scored against, and without it the Service has no way to tell whether the record it found is the right human being. A request missing them is refused with 400 `identity_incomplete` **before any lookup runs and before anything is charged**. Everything else — phone, email, street address, city and state, date of birth, usernames — is optional, and the more of it you send, the more the query planner has to work with: it leads with a held identifier where it has one, in the order **phone** → **email** → **street address** → **name**. A field we cannot parse is dropped rather than guessed at, and the name of the dropped field is returned to you in `dropped_fields`, so a formatting mismatch surfaces as a line in the response instead of a thin result you cannot explain.

Where a call goes — and Verification and Scan are not the same.

- **A Verification reaches exactly one third party: the search provider.** It makes **no model calls**. It reaches no identity-enrichment vendor, no breach or stealer-log vendor, no court-records source, and no username site. One `site:`-scoped query, one provider, one answer. (A failover search provider may stand in for the primary where the primary fails. That is the same single category of recipient and the same single call — not an additional one.)
- **A Scan reaches more, and we will not blur the two together.** In addition to the search provider, a Scan may query **identity-enrichment sources, breach and stealer-log sources, court-record sources, and public username pages**, and it uses **model inference** for parts of its adjudication and phrasing. Some of those sources are **commercial subscription data** — an enrichment vendor and breach-corpus vendors are paid services, not public web pages.

Section 6.1 states what each category of recipient actually receives.

Coverage on a Verification is your list, not ours. There is **no domain allowlist**. You name any hostname and we query it; reach is not narrower than "anything you name" — it *is* anything you name.

The one exclusion. A request naming a domain in the facial-recognition and biometric identification class — a service whose primary function is identifying a person from an image or a biometric template — is **refused with 403** `domain_excluded` **before any query is built and before any money is spent**. That is a control in the code, not a policy we apply by hand, and it is not waivable.

Coverage on a Scan, stated precisely. The base sweep queries a **pinned set of twelve people-search domains**; a typical run issues on the order of thirty to fifty site-scoped queries in total, and the number varies with what you send us and what comes back. The **Registry** — a set of several hundred domains (currently 548) derived from the California data-broker register, a public record you can audit, plus a curated priority set — is what findings are **classified against** and what coverage is **reported over**. **It is not the query list, and a Scan does not query every domain in it.** Every Scan returns a coverage block naming the domains in the expected set that were not reached, up to a per-response limit. We publish what we missed; that is unusual and it is deliberate.

What the Service is not. An Answer is a **point-in-time observation of an index**. It is not a statement about a person. The Service is **not** a consumer reporting agency, credit bureau, background-check, tenant-screening, or employment-screening product; **not** a facial-recognition or biometric system; **not** a law-enforcement tool; **not** a location service; and **not** legal advice. See the AUP and Section 7.

How the Service behaves toward third-party sites. The crawler the Service uses to fulfill your calls identifies itself honestly by user agent, does not log into third-party services, does not solve or bypass captchas, does not circumvent technical access controls, and **honors robots.txt, failing closed** — including when it cannot retrieve or parse a

robots.txt at all. That describes the Service; it is not a representation about Sirveil's other internal systems.

***In plain English.** A check is one question to one search engine about one website, and nothing about it touches an AI model or a data vendor. A sweep is a bigger machine: it also asks enrichment, breach and court sources, some of which we pay for, and it uses models on tokenized data to sort what comes back. Those are two different products and we are not going to describe them with one sentence. Send us a name plus a phone number or an email if you want a result worth paying for. And there is one kind of website we will not check for you, no matter who asks: face-search engines. That refusal happens before we spend a cent.*

6.1 Subprocessors, and what each one receives

To answer your call, the Service sends data to third parties. Stated by category, honestly:

Category	Used by	What it receives
Search / retrieval	Verification and Scan	The site-scoped query, containing the Identifier in clear text
Identity enrichment	Scan only	Identity fields in clear text
Breach and stealer-log lookup	Scan only	The Identifier in clear text; for some sources an email address appears in a request URL
Court records	Scan only	A name-based query
Public username pages	Scan only	Only a handle you supplied, in a profile URL
Model inference	Scan only	Tokenized placeholders only — never a raw name, phone number, email address, street address, date of birth, or government identifier

Retrieval vendors necessarily receive the Identifier in clear text. That is what a lookup is — there is no way to ask an index whether a phone number appears on a page without sending the phone number. We are not going to pretend otherwise, and any vendor who tells you their retrieval works differently is describing a different operation.

The model path is different, and it is enforced rather than promised. Sensitive values are tokenized before any model call, and a check runs on the outbound payload: if a raw sensitive value survives tokenization, the call is **aborted rather than sent**. Model routing is pinned to US-hosted providers, configured to deny data collection and to request zero data retention.

The current list of named subprocessors is available on request at support@sirveil.ai, and we will provide it to a Marketplace or enterprise buyer under a confidentiality undertaking. We do not publish the list, because it is also a map of our supply chain. Sirveil may add, remove, or replace a subprocessor within these categories.

***In plain English.** A check goes to one search provider and stops there. A sweep goes to more places, and the ones that do lookups have to see the real identifier — that is how lookups work. The AI models are the exception: they only ever see PERSON_1 and PHONE_2, and if a real value ever leaked into that path the call gets killed instead of sent. Ask us for the vendor list and you will get it; we just do not hand it to our competitors from a public page.*

6.2 How you call it: shapes, timing, and the agent surface

Three shapes, one implementation.

1. **Synchronous — the default.** The Service returns the result body when the run finishes.

2. **Streaming.** Accept: application/x-ndjson yields progress events followed by the identical result body.
3. **Asynchronous — strictly opt-in, by a named header.** Send Prefer: respond-async and you receive 202 with a job id; poll GET /api/v1/jobs/{id} for the result. Async is **not** the default shape, and an integration built as though it were will receive a result body where it expected a job id.

Async mechanics. A job holds its input and its result for **24 hours**, after which the identity and the result are purged and the job's id, status and timings remain, so a late poll receives an explicit expiry rather than an answer indistinguishable from a bad id. **Polling a job is free and consumes no quota.** A job id belonging to another tenant returns 404, never 403 — we do not confirm the existence of someone else's job. If async is unavailable on a deployment the call returns 503 async_not_available and charges nothing.

Free endpoints. GET /api/v1/whoami (which reports the plan, limits, and channel that apply to your account) and GET /api/v1/jobs/{id} are **free and consume zero quota**.

Timing, as measured — not as promised. In Sirveil's own instrumented runs, a Verification had a **median latency of 890 ms and a 95th-percentile latency of 1,463 ms** (n=18, cache-cold), and a synchronous Scan had a **median of 154 seconds and a maximum of 203 seconds** (n=5). Those are measurements of our runs on our data, not a performance commitment, and Section 15 disclaims any warranty about them. Two practical consequences: do not describe a Verification as "about a second," and **set a read timeout of at least 240 seconds on a synchronous Scan** or use the async shape.

The agent surface. POST /api/v1/mcp is a **read-only** Model Context Protocol server exposing a **single verification tool**. There is no scan tool. A call through it is a Verification and **bills as a Verification**; we say that here because a buyer who discovers it on an invoice is a refund request.

In plain English. A sweep takes about two and a half minutes. That is the product, not an outage — if your HTTP client gives up after thirty seconds, that is a timeout you set. Use Prefer: respond-async if you cannot hold the connection, and poll for free. Checking your key and polling a job never cost anything. If your AI agent calls us over MCP, that is a check and it bills like a check.

7. Permitted use, and the fence

Permitted use is broad. You may use the Service and the Output for any lawful business purpose, including inside products you build and sell, subject only to this Section and the AUP. Inside the fence, your business is your business.

The fence. You will not, and will not permit anyone else to, use the Service or the Output:

1. **for any purpose regulated by the Fair Credit Reporting Act**, 15 U.S.C. §1681 *et seq.*, or a state analogue including Cal. Civ. Code §§1785 and 1786 — including as a factor in establishing a person's eligibility for credit, insurance, employment, housing or tenancy, education, a government license or benefit, or any other purpose enumerated in §1681b(a)(3);
2. **to stalk, harass, dox, intimidate, threaten, or locate a person for the purpose of causing harm**, or to facilitate anyone else doing so;
3. **for facial recognition or biometric identification** — you will not submit face images, face embeddings, or other biometric identifiers or templates, and you will not use Output to build, train, operate, or enrich a facial-recognition or biometric-identification system;
4. **with credentials or sensitive material you have no lawful right to submit** — passwords, password hashes, session tokens, full payment-card data, government-ID images, or data from a breach or dump;
5. **to attack, probe, overload, reverse engineer, or circumvent** the Service, its security, or its metering, or to test it beyond documented limits without our written agreement;

6. **to misrepresent Output** as a government record, a background check, a credit or consumer report, a certification, or a determination about a person; or
7. **in violation of any applicable law**, or of a third party's terms of service, where you direct the Service at that third party.

Your representations about what you send. For every Identifier you submit, you represent that you hold it lawfully, that you have a lawful basis to have us process it, and that you are not knowingly submitting the Identifier of a person you know or reasonably should know is under 18. **The Service does not screen for age and performs no date-of-birth boundary check; this representation is the control, and it is yours.** Nothing in these Terms obliges Sirveil to detect, verify, or refuse an Identifier relating to a minor. **Flow-down.** If you make Output available to your own customers, you will impose these same restrictions on them by contract, require each of them to impose the same restrictions on every further recipient, and name Sirveil as an intended third-party beneficiary entitled to enforce them. You remain responsible for your customers' compliance.

We may check. We may monitor usage patterns for signs of prohibited use, ask you to explain a pattern, and require the certification described in Section 5. Nothing here obliges us to monitor, and no failure to detect misuse waives anything.

***In plain English.** Build whatever you want on top of this — resell it, put it in your own product. The prohibitions are the seven numbered items above: no judging anyone's eligibility for anything (that is background-check territory and we are not a background-check company), no hunting people, no faces or biometrics, no stolen material, no attacking or gaming the meter, no passing our answers off as a background check or a government record, and nothing illegal. If you resell, your customers are bound by the same seven, and so are theirs. One more thing, said plainly: we do not check whether the person you are asking about is a child. You are telling us they are not.*

8. Data protection, roles, and retention

You are the controller of the Identifiers you submit and you decide why they are submitted. For those Identifiers, Sirveil acts as your **service provider** (as defined by the CCPA) and processor: we use them to answer your call and for the limited retained purposes described in the API Privacy Notice, and we do not sell them, share them for cross-context behavioral advertising, retain them beyond the periods stated below, combine them with other data except as permitted, or use them to build or enrich a separate product.

What is kept, and for how long — all three cases.

1. A **synchronous** Verification or Scan **does not persist the identity you submitted or the Answer we returned.** The verification path is guarded in code against writes.
2. An **asynchronous** job must hold its input and its answer between submit and poll. That job is retained for **24 hours**, after which the identity and the result are purged and only the job's id, status and timings remain.
3. Every authenticated call writes **one activity row** — the account, the key, the endpoint, the timestamp, the status, the latency, the units billed, the outcome, an internal run identifier, and our own cost of serving the call. It contains **no request body and no response body**, and therefore no Identifier and no Answer. That row is the billing record, and it is **retained indefinitely**.

We would rather write the third line than tell you we keep nothing, because we do keep that one row and an invoice you can audit depends on it.

A short-lived in-memory query cache also holds normalized query text, partitioned by account, for a period measured in minutes. See Section 10 for how a cache hit bills.

Sirveil is an independent business with respect to your account, billing, and abuse-prevention data about *you*.

Aggregate operational statistics. We may compute and use statistics about how the Service runs — call volumes, latencies, error and refusal rates, capacity and cost — to operate, secure, support, and improve it, and to report on the Service in aggregate. **This right is limited to operational measurements and does not reach the Identifiers you submit or the Output we return.** We do not build a dataset about Subjects, we do not enrich a separate product with your queries, and nothing we publish will identify you or any Subject or make your usage attributable to you.

Service-provider terms. These Terms are the parties' service-provider contract for the purposes of the California Consumer Privacy Act, and no separate addendum is required for Sirveil to act as your service provider. Sirveil:

1. **will not sell or share** the Identifiers you submit, as "sell" and "share" are defined by the CCPA;
2. **will not retain, use, or disclose** them for any purpose other than the business purposes specified in these Terms and the API Privacy Notice — namely, answering the call you made and maintaining the billing record described above — including outside the direct business relationship between us;
3. **will not combine** them with personal information received from any other source, except where the CCPA permits a service provider to do so;
4. **will comply with the obligations the CCPA places on a service provider**, and will provide the same level of privacy protection the CCPA requires of you;
5. **will notify you** if Sirveil determines that it can no longer meet those obligations;
6. **grants you the right** to take reasonable and appropriate steps to confirm that Sirveil uses the Identifiers consistently with your own obligations under the CCPA, and to stop and remediate any unauthorized use — including, on written request, a written description of the measures Sirveil applies; and
7. **acknowledges** that it understands these restrictions and will comply with them.

Where your own regulator or customer requires a separate addendum on their paper, write to support@sirveil.ai and Sirveil will review it.

In plain English. If you call us and wait for the answer, we do not keep the person's details or the result — the answer goes to you and nothing about the subject is stored. If you use the async option, we have to hold the input and the answer until you come back for them, and we delete both after 24 hours. And there is one line per call in our billing ledger — who called, which endpoint, how long it took, what it billed — with nothing about the subject in it. We keep that one. That is the honest version.

9. Honest answers, and their limits

The three-state contract. A Verification returns exactly one of three states — `indexed`, `not_indexed`, `indeterminate` — and the Service ships the limits of the method **inside every response body**, in the `contract.limitation` field, in these words:

"THREE STATES ONLY. This endpoint reads a search index; it does not fetch the page. `not_indexed` means absent from the search index — not proven absent from the site."

That text is reproduced here so that the contract and the payload cannot drift apart. If the two ever differ, the text the Service returns in the response body is the one that describes the Service.

What follows from it, as commitments about how we report:

- `indeterminate` means we could not confirm either way. **It is not a euphemism for `not_indexed`, and we will not report it as one.**
- An identity submitted with a **name and nothing else can never return `indexed`**. A name alone cannot distinguish your subject from everyone who shares the name, so the honest answer is `indeterminate`, and that is the answer you will get.

- `not_indexed` is reserved for the case where a name-shaped query actually ran and returned nothing. Where only one narrow query shape could be tried, the answer is indeterminate rather than a clean negative.
- Public indexes change without notice. An Answer describes what we observed when we observed it, and may be wrong or stale by the time you act on it.
- We do not warrant that Output is accurate, complete, or current, and Output is never a determination about a person's character, conduct, or eligibility.
- The Registry changes as sources appear, disappear, and change behavior.
- Every response — **including every error response** — carries the pipeline, weighting, and calibration version stamps that produced it, so an answer you were given can be traced to the exact logic that produced it.

***In plain English.** This section is the product. When we cannot tell, we say we cannot tell, and we do not dress it up as a clean "no." `not_indexed` means we did not find it in the index we searched — not that it is not on the site. If you send us only a name, expect indeterminate; that is the system being honest, not broken. And every answer, including every error, is stamped with the version of the logic that produced it, so you can always ask us what happened.*

9.1 Confidence, calibration, and linkage are diagnostics, not determinations

A Verification returns confidence and `calibrated_probability`. A Scan returns per-field linkage signals and match probabilities. **They are returned to explain an answer, not to measure its accuracy, and you must not rely on them as a measure of accuracy.**

Stated plainly, because you would otherwise have to take these numbers at face value: the calibration behind `calibrated_probability` is fitted on a small internal sample with very few positive examples, and the linkage weighting is **not calibrated** — the artifact carries that fact on its face. Sirveil has **not** conducted a controlled accuracy study against a customer population and publishes no precision figure for the Service.

Accordingly, you will not (a) present any of these values to a Subject or to any person as a probability or a statement about that person, (b) use one as a threshold or an input in any decision about a person — which Section 7 prohibits in any event — or (c) represent to anyone that a Sirveil number expresses a measured accuracy rate. The parts of an Answer built to be relied on are the **state**, the **evidence**, the ordered **explanation** of what was queried, and the **coverage block**.

***In plain English.** We return our own uncertainty numbers because hiding them would be worse. But do not build a threshold on them, and never put one in front of the person it is about. If you need something you can defend, use the three states, the evidence, and the list of what we actually looked at — those are measured, those are ours, and those we will stand behind.*

9.2 When an upstream source is unavailable

The Service depends on third-party sources. They fail, rate-limit, and change behavior without telling us.

- Where the search provider is unavailable, a Verification returns indeterminate **with an outcome of unserved, and bills you nothing**. We will not charge you for the sentence "we were unable to look."
- Where a Scan lane's source is unavailable, the lane is reported in the coverage block as failed or never queried, **by name**, rather than silently dropped.
- The unavailability, degradation, or behavior change of a third-party source, and any refusal or indeterminate that results, **is an operating condition of a retrieval service and is not a breach of these Terms**. Section 15 disclaims a warranty of uninterrupted service and Section 24 covers force majeure.

***In plain English.** When somebody upstream breaks, you get an honest "we could not tell" and a zero on the invoice, and a sweep tells you by name which sources it missed. That is the deal working, not the deal failing — so it is not a breach we owe you damages for. What you do get is the truth about what we reached.*

10. Fees, metering, and billing

What you pay. Per-call prices are set out in the Rate Card. The Rate Card is not reproduced here so that a price change does not require an amendment; prices in effect when a call is made are the prices for that call.

What bills, and what does not.

Event	Billed?
A Billable Answer, computed fresh	Yes
A Billable Answer served from cache	Yes — a delivered answer is a delivered answer
An Answer of <code>indeterminate</code> where the index genuinely could not confirm	Yes
A completed call where you dropped the connection before reading the response	Yes
A call through the MCP surface	Yes — as a Verification
An Unservd Call — our fault or an upstream provider's fault	No — zero units
A call refused before any lookup (excluded domain class, bad request, bad key)	No
A call refused because you exhausted a cap or limit	No
<code>whoami</code> , and polling a job	No — free, and zero quota

Cache hits bill in full, in the specification's own words: *"A cache hit is billed and consumes quota exactly like a fresh call. The cache makes a repeat answer consistent, not free."* The cache is per-instance and short-lived — currently minutes, not hours — and we may change how it works. What we will not change without notice is the billing rule in this paragraph, which is stated here and in the Documentation rather than discovered on an invoice.

How billing runs. The Service is sold **metered**. You pay per call, at published rates, for calls that were served. There is **no subscription, no minimum, no seat count, no committed spend, and no free trial**, and Sirveil publishes no plan ladder for the Service. The Service does contain quota mechanics — an account may be provisioned with a monthly unit budget, and a call that would exceed it is refused with `429 quota_exceeded` and bills nothing — but an account purchased through a Marketplace is provisioned **without** a unit budget. Whatever limits apply to your account are reported to you, free, by `GET /api/v1/whoami`.

Price changes. We may change the Rate Card on **at least 30 days' notice** to your account email. Changes apply prospectively only; they never bill retroactively for usage already incurred. Where you purchase through a Marketplace, a price change on that listing also follows the Marketplace's own change and notice process, and the price published on the listing when a call is made is the price for that call.

Metering is the record. Every authenticated call writes exactly one metering row from the Service's own accounting path. That row — not a client-side count — is the billing record. We will show you your usage, and we will investigate any discrepancy you report **within 90 days of an invoice** in good faith and correct genuine errors.

Taxes. Prices exclude taxes. You are responsible for all taxes, duties, levies, and similar charges arising from your purchase other than taxes on Sirveil's net income. Payments are made without deduction or withholding; if withholding is required by law, you will gross up so that Sirveil receives the amount it would have received without the withholding. Where you purchase through a Marketplace, that Marketplace's tax handling applies to the transaction.

If an invoice goes unpaid. An undisputed amount not paid when due accrues interest at the lesser of **1% per month** or the maximum the law allows, from the due date until paid, and you will reimburse reasonable costs of collection, including legal fees. We will tell you before we suspend, and Section 18 gives ten days' notice. **An amount you dispute in good faith under the 90-day process above does not accrue interest while it is being investigated.**

Marketplace billing. If you purchased through a Marketplace, invoicing, payment, refunds, and disputes follow that Marketplace's process, on usage Sirveil reports to it. Your invoice line items read `verification` and `scan`.

***In plain English.** You pay for answers we delivered. A cached answer took work to build and is still an answer, so it bills — we tell you that up front rather than hiding it. If our side broke, you pay nothing. If you hung up before reading a completed scan, that one is on you: use the `async` option if your client cannot wait, and the documented timeout is there so this never surprises you. No monthly fee, no minimum, no trial to forget to cancel — you pay for calls and nothing else.*

10.1 What bounds your bill — and what the spend ceiling actually caps

Your bill is bounded by **your own call volume**, and by a unit quota if your account has been provisioned with one. Nothing else bounds it.

In particular, the per-account protective ceiling described in Section 11.1 is measured against **Sirveil's own supplier costs, not against your charges**. It exists to stop a runaway integration from costing us money. **It is not invoice protection, it does not guarantee your bill stays under any figure, and we will never tell you that it does.** If you need a hard spending bound, bound the number of calls your integration is able to make — that is the only control that actually caps a metered bill, and it is on your side of the line.

***In plain English.** Nobody can promise you a small bill except you. Our ceiling protects our costs, not your invoice, and a vendor who sells you their cost cap as your spend cap is selling you something that is not true. Count your calls.*

10.2 No minimum, no volume commitment, no reserved capacity, no exclusivity

- You commit to **no volume**. If you make no calls, you owe nothing.
- Sirveil **reserves no capacity** for you and holds no headroom for your traffic. Capacity is shared and subject to Section 11.1. If you need reserved capacity, that is an enterprise agreement, not this one.
- Nothing here is **exclusive** in either direction. We may serve anyone, including your competitors. You may buy the same capability from anyone else, including while you buy it from us.

11. Rate limits, changes to the Service

We may set and adjust rate limits, concurrency limits, and quotas, and may throttle or queue calls to protect the Service and other customers. Documented limits will be published, and the limits applying to your account are readable at any time, free, from `GET /api/v1/whoami`. Limits are applied on a best-efforts basis: serving a call that exceeded a published limit does not raise your limit, waive it, or entitle you to that behavior again.

We may add, change, or discontinue features; where a change would materially reduce the Service you are actively paying for, we will give at least 30 days' notice and you may stop using the Service before it takes effect.

Support and availability. No uptime commitment applies to self-serve accounts, and none is offered here. Sirveil publishes no status page or availability metric for the Service, and any monitoring, status, or incident information we may publish or share is provided for information only and is not a service-level commitment. Service-level commitments are available only under a signed enterprise agreement.

11.1 Protective limits, refusals, and what they mean

Four independent bounds can refuse a call. They are checked in a deliberate order so you always receive the most specific and most actionable refusal available:

Bound	Refusal	Billed
Rate limit — calls per minute for your account	429 rate_limited, with Retry-After	No
Unit quota — where your account has one	429 quota_exceeded	No
Per-account protective ceiling — Sirveil's supplier spend for your account	402 spend_ceiling_reached	No
Platform-wide daily capacity guard	503 capacity_reached	No

A suspended account is refused 403 tenant_suspended. A request naming an excluded domain class is refused 403 domain_excluded. A malformed or incomplete request is refused 400. An unavailable retrieval source returns 503 search_unavailable or, on a Verification, an indeterminate answer with an unserved outcome.

A refusal under any of these is a documented limit operating as intended. It is not a breach of these Terms, and it does not entitle you to damages, a credit, or a refund. None of them bills you anything. Where a limit refuses you and you think the limit is wrong for your account, write to support@sirveil.ai — that is a conversation, and for most accounts it is a configuration change rather than an argument.

In plain English. Four things can say no: you are calling too fast, you used up a quota, your account hit our internal cost guard, or the whole platform hit its daily guard. All four are free, all four tell you which one it was, and none of them is us breaking the contract. If one of them is getting in your way, email support — we would rather raise a number than lose a customer over one.

11.2 Features offered for evaluation

We may label a feature, endpoint, or response field **beta, preview, experimental, or evaluation**. Anything so labelled is offered **as is, for evaluation only**, may change or be withdrawn **without the notice Section 11 otherwise requires**, is excluded from the commitments in Section 15, and should not be relied on in production. Where such a feature bills, it bills at the published rate and the metering warranty in Section 15(a) still applies to it — **we will not charge you for a call we did not serve, whatever the feature is called**.

In plain English. If we mark something beta, treat it as beta: it may move or vanish, and we are not promising it will keep working. The one thing that does not change is the meter.

12. Your license to Output — deliberately broad

Subject to Section 7, to the limits in this Section, and to payment, Sirveil grants you a **perpetual, worldwide, non-exclusive, royalty-free license** to store, reproduce, modify, create derivative works from, embed, white-label, bundle, display, distribute, and **resell** the Output, alone or within your own products, at any price you choose, with no revenue share and no additional commercial fee.

Two limits, and they cost an honest customer nothing.

(a) Not for competing use. This license does not extend to using Output to build, train, or enrich a dataset, model, or service that detects, scores, or reports personal-data exposure in competition with the Service, or to replicating the Registry. Sell the answers; do not assemble them into the thing that answers.

(b) Conditioned on the fence. This license is conditioned on your compliance with Section 7 and AUP Section 4. If we give you written notice of a breach of either and you do not cure it within 15 days, this license ends **as to Output delivered to you after the date of that notice**, whether or not we also terminate the agreement. Section 19 governs Output already delivered where we terminate for a fence breach.

Output may incorporate material sourced from third parties. Sirveil grants only such rights as it holds. We do not represent that reproducing or distributing evidence drawn from a third-party page is permitted by that party, and clearing third-party rights in such material is your responsibility.

This license is limited by Sections 7 and 13 and by the AUP, and by the flow-down obligation in Section 7; use outside those limits is outside the scope of this license. Subject to Section 19, it survives termination as to Output already delivered. Sirveil claims no ownership of your products or of your customers' relationships.

***In plain English.** What you paid for is yours, and it stays yours after you stop being a customer — unless we had to cut you off for breaking the fence, which is the one case where the resale right stops (Section 19). Build a business on it. Two strings, both narrow: do not use our answers to build a rival to the thing that produced them, and keep to the seven prohibitions in Section 7 — if you break those and do not fix it, new answers stop being licensed. Everything else travels with the data to whoever you sell it to, and to whoever they sell it to.*

13. Our intellectual property

The Service, the Registry, the Documentation, the software, the Sirveil marks, and everything we make available other than Output are owned by Sirveil and licensed to you only as set out here. You get no rights by implication or estoppel; all rights not expressly granted are reserved.

You will not use the Service, the Documentation, or Output to:

1. build, train, or enrich a dataset, model, or service that competes with the Service;
2. replicate the Registry; or
3. conduct benchmarking or competitive analysis **for, or on behalf of**, a person who offers or is developing a service that competes with the Service.

What this does not restrict. You may evaluate the Service for your own purposes, run your own tests against it, publish your own honest results, and say what you found. We publish comparisons ourselves and we are not going to forbid you from publishing yours. What is out of bounds is doing that work as a competitor's proxy, and copying the thing rather than using it.

Feedback. If you send us ideas or suggestions, we may use them freely, without obligation or compensation, and they are not your confidential information.

Publicity. Neither party will use the other's name or marks publicly without consent, except that Sirveil may identify you as a customer by name and logo in customer lists and on its website; tell us at support@sirveil.ai if you would rather we did not, and we will stop.

14. Confidentiality

Each party will protect the other's non-public information disclosed in connection with the Service — including the Rate Card where individually negotiated, security details, subprocessor lists disclosed under Section 6.1, and unreleased roadmap — with at least reasonable care, use it only for the Service, and not disclose it except to personnel and advisers bound to comparable terms. The usual exclusions apply (already known, public through no fault, independently developed, lawfully received). Compelled disclosure is permitted with prompt notice where lawful. Obligations run for three years after disclosure, and indefinitely for trade secrets.

15. Warranties and disclaimer

What we commit to — three specific things, and no general one.

(a) We warrant that we meter honestly. A call is billed only where the Service actually served it. This is not a policy we apply by hand: the metering path computes billable units as *the units requested if, and only if, the outcome was served, and zero otherwise*, so a non-2xx response, a refusal, a provider failure, and an unserved answer all carry zero billable units by construction rather than by intention. If we ever bill you for a call we did not serve, that is a breach of this warranty and we will correct it.

(b) Published limits and named refusals. The limits that apply to your account are published in the Documentation and readable free from `GET /api/v1/whoami`, and a refused call returns a named refusal code from the published list rather than a generic error.

(c) Notice before a material reduction. We will give the notice described in Sections 11 and 20 before a change that would materially reduce the Service you are actively paying for.

We give no general warranty of skill, care, workmanship, availability, or performance, and none is to be implied. We have not undertaken a service level, and Section 11 states that no uptime commitment is offered.

You warrant that you hold the Identifiers you submit lawfully, that your use complies with Section 7 and the AUP, and that you have made the disclosures and obtained any consents your own use requires.

Otherwise, the Service and the Output are provided "AS IS" and "AS AVAILABLE." TO THE MAXIMUM EXTENT PERMITTED BY LAW, SIRVEIL DISCLAIMS ALL OTHER WARRANTIES, EXPRESS OR IMPLIED, INCLUDING MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, ACCURACY, AND NON-INFRINGEMENT, AND DOES NOT WARRANT THAT THE SERVICE WILL BE UNINTERRUPTED OR ERROR-FREE OR THAT OUTPUT IS ACCURATE, COMPLETE, OR CURRENT. **THIS DISCLAIMER EXTENDS TO ANY WARRANTY ARISING FROM COURSE OF DEALING, COURSE OF PERFORMANCE, OR USAGE OF TRADE**, AND TO ANY GUARANTEE, PROMISE, ASSURANCE, OR UNDERTAKING NOT WRITTEN IN SECTION 15(a) TO (c) ABOVE. **NO PERSON — INCLUDING ANY EMPLOYEE, AGENT, OR RESELLER OF SIRVEIL — IS AUTHORIZED TO GIVE ONE**, AND ANY PURPORTED GUARANTEE OR PROMISE GIVEN OUTSIDE THESE TERMS IS OF NO EFFECT.

Statements on our website and in our marketing about our mission, our roadmap, or the future of personal-data control are made in good faith, are aspirational, and are not warranties or part of this agreement. Latency, throughput, and accuracy figures we publish are measurements of our own runs and are not commitments.

***In plain English.** We promise one thing hard: we will not bill you for a call we did not serve, and that is enforced in the code rather than trusted to us. We also promise to publish the limits and to warn you before we take something away you are paying for. What we do not promise is that the service is always up, always fast, or always right. We have not signed up to a service level, and we are not going to imply one in softer words: a promise nobody is measuring is a promise we would be asking you to take on faith.*

15.1 No certifications, no audit, no penetration test

Stated expressly, so that nothing is implied and so you learn it now rather than in diligence:

Sirveil holds no SOC 2 report, no ISO 27001 certification, and has not had a third-party penetration test. Sirveil holds no other security or privacy certification, attestation, or audit report for the Service, and nothing in these Terms, the Documentation, or our marketing should be read as claiming one.

Data is encrypted in transit and is subject to our platform providers' encryption at rest. Beyond the specific commitments in these Terms and the API Privacy Notice, we make no representation about particular security controls.

If your organization requires a certification, an audit report, or a completed penetration test as a condition of purchase, the answer today is that we do not have one. Ask before you subscribe; we would rather lose the deal in week one than fail your review in month six.

***In plain English.** No SOC 2. No ISO 27001. No pen test. If you need those, we are not ready for you yet, and we would rather say so on this page than in a security questionnaire six months from now.*

16. Indemnities

You indemnify us. You will defend, indemnify, and hold harmless Sirveil and its officers, directors, employees, and agents from third-party claims, and resulting damages, losses, and reasonable legal costs, arising from (a) your or your customers' use of the Service or Output in breach of Section 7 or the AUP, (b) Identifiers you submitted without a lawful right to do so, (c) your breach of these Terms or of applicable law, or (d) claims by a Subject arising from your use of Output. We may assume control of the defense at your expense, and you will cooperate.

We indemnify you, and here is the size of it. We will defend you against a third-party claim that the Service, used as documented and within Section 7, infringes that party's US intellectual property, and pay damages finally awarded or amounts we agree in settlement. This does not apply to claims arising from Output content, from Identifiers you supplied, from your modifications or combinations, or from use outside the Documentation or Section 7. If the Service becomes the subject of such a claim, our remedies — and yours — are that we may procure a right to continue, modify the Service, or terminate access and refund prepaid unused fees. **Our defense costs, together with any damages or settlement amounts we pay under this paragraph, count against the cap in Section 17(b), and this paragraph is our entire liability for infringement claims.**

***In plain English.** If you misuse the Service and someone sues us for it, that is yours to defend. If someone claims our Service infringes their IP, we will defend you — up to the same ceiling that applies to everything else we owe, with the defense costs counted inside it, and with our options being to fix it, license it, or shut it off and refund you. We are not writing a blank check for a lawsuit; we are telling you exactly what the check is.*

17. Limitation of liability

TO THE MAXIMUM EXTENT PERMITTED BY LAW:

(a) NEITHER PARTY IS LIABLE FOR INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL, OR PUNITIVE DAMAGES, OR FOR LOST PROFITS, LOST REVENUE, LOST DATA, OR LOSS OF GOODWILL, EVEN IF ADVISED OF THE POSSIBILITY.

(b) EACH PARTY'S TOTAL AGGREGATE LIABILITY ARISING OUT OF OR RELATING TO THESE TERMS OR THE SERVICE IS LIMITED TO THE GREATER OF (i) THE FEES YOU PAID SIRVEIL FOR THE SERVICE IN THE TWELVE MONTHS BEFORE THE EVENT GIVING RISE TO THE CLAIM, OR (ii) ONE HUNDRED US DOLLARS.

(c) Neither the exclusions in (a) nor the cap in (b) applies to: your payment obligations; your indemnity in Section 16; your breach of Section 7 or the AUP; **your** breach of Section 14; or either party's fraud or willful misconduct. Sirveil's liability for its own breach of Section 14 remains subject to the cap in (b).

(d) These limits apply regardless of the theory of liability and survive any failure of essential purpose. Nothing here excludes liability that cannot lawfully be excluded.

***In plain English.** Neither of us can be bankrupted by the other over a call that costs less than a dollar. The ceiling is roughly what you paid us in the past year. Paragraph (c) lists the situations where the ceiling comes off — including one that runs against us, if we act fraudulently or with willful misconduct.*

18. Suspension

We may suspend a key, an account, or a Marketplace entitlement immediately, without liability, where we reasonably believe there is a breach of Section 7 or the AUP, a security or fraud risk, a legal requirement, or a threat to the Service or to others; and on **ten days' notice** where an undisputed invoice is unpaid. Where practicable and lawful we will tell you why and give you a chance to fix it first; where the risk is immediate we may act first and tell you after. A suspended account is refused 403 `tenant_suspended`. Suspension does not relieve you of fees for Billable Answers already delivered.

19. Term, termination, survival

These Terms run while you have an account or continue to call the API. You may terminate at any time by closing your account and stopping calls; there is no exit process and no termination fee. We may terminate for material breach not cured within 15 days of notice, immediately for a breach of Section 7 or of AUP Section 4, on 15 days' notice to cure for other AUP breaches, and otherwise on 30 days' notice.

On termination: your right to call the API ends; fees already incurred remain payable; the license in Section 12 continues for Output already delivered; and your data is handled as Section 8 and the API Privacy Notice describe — a pending asynchronous job expires on its ordinary 24-hour clock, and the per-call activity rows described in Section 8(3) are retained as the billing record.

Except where we terminate for breach of the fence. If we terminate for your breach of Section 7 or of AUP Section 4, the license in Section 12 ends as to further reproduction, distribution, and resale; you will stop distributing Output and delete it on request. You may continue using Output already embedded in a delivered product only so far as necessary to avoid harm to a third party, and not to make new copies.

Sections 2, 3, 7 (as to Output retained), 8, 9 (with 9.1 and 9.2), 10 (as to accrued fees, with 10.1 and 10.2), 11.2, 12–17 (with 15.1), and 19–24 survive, together with any provision that by its nature should survive.

20. Changes to these Terms

We may update these Terms on at least **30 days' notice** to your account email. Changes apply prospectively; they never bill retroactively and never narrow the license in Section 12 for Output already delivered. If you do not accept a change, stop using the Service before it takes effect. Continued use after the effective date is acceptance. Each version carries a version number and effective date.

21. Compliance

You will comply with applicable export-control, sanctions, and anti-corruption laws, will not export or re-export the Service or Output in violation of them, and will not use the Service on behalf of a restricted party.

US Government end users. The Service and the Documentation are "**commercial products**" comprising "**commercial computer software**" and "**commercial computer software documentation**" as those terms are used in **48 C.F.R. § 2.101**. If acquired by or on behalf of the US Government, they are provided with **only those rights granted to all other end users under these Terms**, in accordance with **48 C.F.R. § 12.212** and **48 C.F.R. §§ 227.7202-1 through 227.7202-4**. No other licence, right, or government data right is granted, and any provision inconsistent with this paragraph does not apply.

22. Disputes — talk to us first, then arbitrate

Step one, mandatory for both sides. Before starting arbitration or any formal proceeding — other than an application for injunctive or emergency relief permitted by the carve-outs below, or a small-claims action — the complaining party will send a written **remedy petition** — to support@sirveil.ai for us, to your account email for you — describing the dispute and the remedy sought. The parties then have **60 days** to resolve it in good faith, including at least one call or video conference between people with authority to settle. Limitation periods are tolled for those 60 days.

Step two. Any dispute not resolved that way is finally resolved by **binding arbitration** administered by the American Arbitration Association under its Commercial Arbitration Rules, before one arbitrator, in English, seated in **San Jose, California**, with remote hearings available. Judgment may be entered in any court of competent jurisdiction. The Federal Arbitration Act governs this Section.

Individual basis only. CLAIMS ARE ARBITRATED INDIVIDUALLY; NO CLASS, COLLECTIVE, CONSOLIDATED, OR REPRESENTATIVE PROCEEDINGS, AND THE ARBITRATOR MAY NOT CONSOLIDATE CLAIMS.

Bench determination for claims that reach a court. For any claim within the carve-outs below, or any claim otherwise heard by a court: to the extent a pre-dispute waiver of jury trial is enforceable under the governing law, both parties waive trial by jury. **Because California does not enforce pre-dispute jury-trial waivers, the parties instead agree that any such claim heard in a California court will be submitted to a referee for a general reference under California Code of Civil Procedure §638(a), the referee to try all issues, whether of fact or of law, and to report a statement of decision on which judgment may be entered.** The parties will agree on a referee within 15 days of a written request; failing agreement, the court appoints one. The parties share the referee's fees equally, subject to any award of costs in the final decision. If the general-reference provision is held unenforceable, the claim is tried by the court, and the jury waiver applies only where enforceable.

Carve-outs. Either party may bring a qualifying claim in small-claims court, and either party may seek injunctive relief in court for breach of Section 7, intellectual-property misuse, breach of confidentiality, or security abuse. Emergency and injunctive relief is not subject to the 60-day waiting period in step one.

Time limit. Any claim must be brought within **one year** after it accrues or it is permanently barred, to the extent the law permits. This does not apply to claims for unpaid fees, to Sirveil's claims under Sections 7, 13, or 16, or to a claim that could not reasonably have been discovered within that period — fence breaches and registry misuse routinely surface late, through a regulator or a downstream customer.

23. Governing law

California law governs these Terms, without regard to conflicts rules. Subject to Section 22, the state and federal courts in Santa Clara County, California have exclusive jurisdiction. The UN Convention on Contracts for the International Sale of Goods does not apply. The choice of California here is deliberate and is not an oversight against Sirveil's consumer Terms of Service, which specify Delaware. These are different products with different counterparties, and California is the correct law for this one: Sirveil's principal office and the arbitration seat are both in San Jose, Sirveil has been qualified to transact business in California since 16 September 2025, and the general-reference procedure in Section 22 is a California procedure that assumes California venue.

24. General

- **Assignment.** You may not assign these Terms without our written consent; we may assign in connection with a merger, acquisition, or sale of assets. Any other attempted assignment is void.
- **Force majeure.** Neither party is liable for delay or failure caused by events beyond its reasonable control, including upstream provider outages, network failures, government action, and natural disasters. Payment obligations are not excused.

- **Independent contractors.** No partnership, joint venture, agency, or employment is created.
- **No third-party beneficiaries,** except that Sirveil's officers, directors, employees, and agents may enforce Sections 16 and 17; a Marketplace may enforce the provisions required by its own terms; and Sirveil may enforce the Section 7 flow-down restrictions against your customers and their recipients as an intended third-party beneficiary of those contracts.
- **Notices.** To you, by email to your account address, deemed given when sent. **To us, every notice of any kind goes to** support@sirveil.ai — one monitored channel, deemed given one business day after sending. Put the matter in the subject line so it reaches the right person on the first read: **Legal Notice · Privacy · Security · Abuse Report · Remedy Petition.** A notice is effective whether or not the subject line is right.
- **Severability.** An unenforceable provision is reformed to the minimum extent necessary, or severed, and the rest stands.
- **No waiver.** A delay in enforcing is not a waiver, and no waiver is effective unless written.
- **Entire agreement.** These Terms with the documents in Section 3 are the whole agreement about the API and supersede every prior or contemporaneous agreement, proposal, quote, demonstration, roadmap, benchmark, blog post, social post, marketing page, sales conversation, and AI-assistant output, whether written or oral, on the same subject.
- **No reliance on anything outside this agreement.** You acknowledge that in deciding to use the Service you have relied **only** on these Terms and the documents named in Section 3, and on nothing else — including any statement on sirveil.ai, on any marketplace listing, in any specification, benchmark, sample report, calculator, or illustration, or made by any person on Sirveil's behalf. Those materials describe the Service in general terms, are produced for a general audience, and may be incomplete, out of date, or wrong. **Where any of them differs from these Terms, these Terms govern,** and no such material creates a warranty, a condition, a representation, a guarantee, or a promise of any kind. This paragraph does not limit liability for fraud or for any liability that cannot lawfully be limited.
- **Feedback.** If you send us suggestions, feature requests, or comments about the Service, you grant us a perpetual, worldwide, royalty-free, irrevocable license to use them without restriction or obligation to you. We will not identify you as the source without your consent.
- **Publicity.** Neither party will use the other's name, logo, or marks in publicity without prior written consent, except that you may state factually that you use the Service.
- **Counterparts and acceptance.** These Terms may be accepted electronically, and acceptance by subscribing through a Marketplace or by issuing or using a key is as effective as a signature.
- **Interpretation.** Headings are convenience only; "including" means "including without limitation"; these Terms are not construed against either party as drafter. The plain-English notes are explanatory only, are not operative, do not enumerate exhaustively, and where a note and the operative text differ, the operative text controls.

25. Contact

Sirveil, Inc., a Delaware corporation · San Jose, California General: support@sirveil.ai · Privacy: support@sirveil.ai · Security: support@sirveil.ai · Support: support@sirveil.ai · Remedy petitions: support@sirveil.ai

26. Document control

Field	Value
Version	2.0

Field	Value
Date settled	2026-08-25
Supersedes	The versions of 2026-08-21 and 2026-08-19, and the page previously published at sirveil.ai/scan-api/terms
Basis	Technical deep dive 2026-08-24 (code landing/v34 @ 5ca896e ; live OpenAPI 1.4.0; production data plane queried 24 Aug 2026) · Alignment brief, API legal set v2 · Input Requirements Brief 2026-08-19 · Status of Record 2026-08-20
Effective date	25 August 2026

Drafting principle applied throughout this version. Describe generously; warrant narrowly. Every convenience the Service delivers mechanically — unserved calls billing nothing, refusals billing nothing, free whoami and polling, no allowlist, the method's limits printed inside every response — is *stated*, because a customer should be able to read these Terms and see how favorable the billing is. Each is stated as a **description of how the Service operates**, not as a warranty of how it will always operate. The single exception is honest metering in Section 15(a), which is structurally enforced in the metering path and cannot silently regress, and which is therefore warranted outright.

Added 25 August 2026 — protective provisions that were absent. Non-reliance (§ 24) · warranty disclaimer extended to course of dealing, usage of trade, and to any guarantee or promise given by any person outside these Terms (§ 15) · feedback licence, publicity, counterparts and electronic acceptance (§ 24) · features offered for evaluation (§ 11.2) · aggregate operational statistics, expressly excluding Identifiers and Output (§ 8) · interest and collection costs on unpaid invoices (§ 10) · suspension on ten days' notice for non-payment (§ 18) · US Government restricted rights (§ 21). **All contact is now one channel** — support@sirveil.ai, with **subject-line routing** (§ 24).

Cross-references. Section numbering in this version is unchanged from the previous one — every new section is a decimal suffix — so cross-references from the Acceptable Use Policy and the API Privacy Notice to Sections 3, 7, 14 and 22 still land.

ACCEPTABLE USE POLICY — SIRVEIL EXPOSURE API

About the plain-English notes. Under most sections there is a short note explaining what that section means. The notes exist to help; they are not legal advice, they do not replace the formal text, and where a note and the legal text differ, **the legal text controls**.

1. Why this exists

Sirveil sells a question-answering service about information that is already public. That capability is useful to auditors, security teams, privacy operations, incident responders and counsel — and the same capability would be useful to a stalker. This Policy is the contractual half of controls that already exist in code.

The engineering controls came first. Facial-recognition domains are refused at the API boundary before a query is built. The one part of the Service that fetches third-party pages honors robots.txt and fails closed. Nothing in the product logs in anywhere, solves a captcha, or steps around an access control. This document does not create those behaviors; it makes the rules around them enforceable, and it is careful to tell you which of them we **describe** and which of them we **promise** — those are different things, and Section 6 says which is which.

Read Section 4 first. It is the whole of the fence, and it is the same fence as Section 7 of the Terms — this Policy elaborates it with examples, and it does not add to it.

Breach of Section 4 of this Policy is a breach of the Terms, and it carries three distinct consequences, each with its own mechanic:

- it permits **immediate termination without a cure period** (Terms Section 19);
- on **written notice and a failure to cure within 15 days**, it ends your license to Output **delivered after that notice**, whether or not we also terminate (Terms Section 12(b)); and
- where we **do** terminate for it, your right to reproduce, distribute and resell Output **already delivered** ends as well (Terms Section 19).

It is also one of the few things that **lifts the liability cap** (Terms Section 17(c)). None of that is a threat. It is the reason a customer who stays inside the fence gets terms this permissive.

2. Scope

This Policy applies to:

- every call to the **Sirveil Exposure API**, through every endpoint — POST /api/v1/verify, POST /api/v1/scan, GET /api/v1/jobs/{id}, GET /api/v1/whoami, and POST /api/v1/mcp;
- **every call made through the agent surface**. POST /api/v1/mcp is a read-only Model Context Protocol server exposing a single verification tool. A call an AI agent makes on your behalf is **your** call, made under your key, subject to this Policy in full. An agent is not a third party for these purposes and "the model decided to" is not a defense;
- every use of the **Output**, including by your personnel, your contractors, your end users reaching the Service through your application, and your own customers if you resell; and
- however access is purchased. Section 3.1 of the Terms states which channels can actually transact today; this Policy applies through any of them, and through any reseller arrangement Sirveil enters.

Where this Policy and the Terms differ, **the Terms control** (Terms Section 3).

3. Permitted use — and it is deliberately broad

Any lawful business or professional purpose, subject to Section 4 of this Policy and to your obligations under the Terms. That includes building the Output into your own products, white-labeling it, bundling it, and reselling it at any price you choose, with no revenue share — provided the restrictions in Section 4 travel downstream by contract under Section 7 below.

Inside the fence, your business is your business. We do not ask what your product does, we do not require you to justify a query, and we do not maintain an approved-use list. There is no allowlist of domains you may verify against: you name any hostname and we query it.

***In plain English.** This is a short list of things you must not do, wrapped around a very large space of things you may. We are not going to audit your roadmap. Stay inside Section 4 and build whatever you want.*

4. Prohibited use — the fence, in seven items

This Section is the fence. There is no enterprise override, no research exception, no government exception, and no price at which any of it unlocks.

The seven items below are the same seven prohibitions as **Section 7 of the Terms**, restated with examples so that a compliance reviewer can see what falls inside each one. **The examples illustrate; they do not enlarge.** Nothing in this Section prohibits conduct that Section 7 of the Terms permits, and if a reading of an example here would extend the fence beyond Section 7, Section 7 controls.

You will not, and will not permit anyone else to, use the Service or the Output:

4.1 For any purpose regulated by the Fair Credit Reporting Act

— 15 U.S.C. §1681 *et seq.*, or a state analogue including Cal. Civ. Code §§1785 and 1786 — including as a factor in establishing a person's eligibility for credit, insurance, employment, housing or tenancy, education, a government license or benefit, or any other purpose enumerated in §1681b(a)(3).

Inside this prohibition, by example: pre-employment or contractor vetting; hiring, firing, promotion, reassignment, or retention decisions; gig-platform or marketplace onboarding decisions about a worker or seller; tenant screening or a rental application; insurance underwriting or claims eligibility; a credit or lending decision; an admissions decision; a professional-licensure or government-benefit determination; and any adverse action taken about a person in whole or in part on an Answer.

Not inside it, by example: checking your own organization's exposure; checking the exposure of your own workforce as a privacy-remediation service for them, where no decision about them turns on the result; and a Subject-initiated check about that Subject's own data.

Sirveil is **not** a consumer reporting agency and the Output is **not** a consumer report. Using it as one puts *you* in the position of a user of consumer reports with none of the protections the FCRA requires, and it breaches this Policy.

4.2 To stalk, harass, dox, intimidate, threaten, or locate a person for the purpose of causing harm

— or to facilitate anyone else doing so.

Inside this prohibition, by example: assembling a location or movement profile of a named individual for a requester who will not say why; publishing a person's home address, phone number, or workplace to expose them; running lookups for a client you know is subject to a restraining or protective order concerning the Subject; building or selling a product marketed to find a former partner, an anonymous critic, or a person who has cut off contact; supplying Output to a person you know intends any of the above.

The qualifier is real. The words *"for the purpose of causing harm"* attach to *locate*. Lawful service of process, lawful debt collection, lawful investigative work performed under a license where one is required, and lawful journalism are not prohibited by this item. Stalking, harassment, doxxing, intimidation and threats are prohibited regardless of purpose.

4.3 For facial recognition or biometric identification

You will not submit face images, face embeddings, or other biometric identifiers or templates, and you will not use Output to build, train, operate, or enrich a facial-recognition or biometric-identification system.

The class is defined by **function** — a service whose primary function is identifying a person from an image or a biometric template — not by a fixed list of companies, so that it does not go stale and so that a contract does not name businesses in it. The current list of the **domain classes the Service refuses** is available on request, subject to the confidentiality obligations in Section 14 of the Terms. (The Registry itself — the full source map — is Sirveil's confidential property and is not disclosed.)

You will not attempt to route around that refusal, by aliasing a hostname, by using a redirector or mirror, by splitting a query across accounts, or otherwise.

4.4 With credentials or sensitive material you have no lawful right to submit

— passwords, password hashes, session tokens, full payment-card data, government-ID images, or data from a breach or dump.

Inside this prohibition, by example: pasting a credential dump into an identity field; submitting a stolen or purchased combo list; sending a scan of a driver's license or passport; supplying a full primary account number in any field.

What this does not stop. Asking whether an email address you hold lawfully appears in a known breach corpus is a lane the Service runs. ****That lane reports *whether* an identifier appears — not the credential itself.**** Nothing in the product uses, returns, replays, or tests a breached credential. What this item stops is you sending us the dump.

You will also not use the Service to log into a third-party site, to solve or bypass a captcha, to circumvent a technical access control, or to retrieve material in violation of the Computer Fraud and Abuse Act or a state equivalent. *(That is item 7 of the fence — unlawful use, and use in violation of a third party's terms where you direct the Service there — applied to these facts. It is an example, not a further prohibition, and the Service has no capability to do any of it in any event: see Section 6.)*

4.5 To attack, probe, overload, reverse engineer, or circumvent the Service

— its security, or its metering, or to test it beyond documented limits without our written agreement.

Inside this prohibition, by example: load testing, stress testing, fuzzing, or penetration testing without prior written agreement; probing or tampering with metering, quota accounting, or rate limiting; opening additional accounts, rotating keys, or splitting traffic across entities to evade a limit; crafting requests to recover an internal threshold, weighting, or rubric; attempting to reconstruct Sirveil's supplier set or unit economics from responses; and attempting to reach another tenant's job, key, or usage.

What this does not stop. You may evaluate the Service for your own purposes, run your own tests against it within published limits, publish your own honest results, and say what you found. Section 13 of the Terms says so expressly. What is out of bounds is doing that work as a competitor's proxy, and copying the thing rather than using it.

If you want to run a load test, ask at support@sirveil.ai. The answer is usually yes, with a window and a limit.

4.6 To misrepresent Output

— as a government record, a background check, a credit or consumer report, a certification, or a determination about a person.

Inside this prohibition, by example: labeling a report "background check" or "verified against government records"; presenting confidence, calibrated_probability, linkage, or a match_probability to anyone as a measured accuracy rate or as a statement about a person — Section 9.1 of the Terms is explicit that these are diagnostics and not determinations; reporting an indeterminate result to a downstream recipient as a clean negative; and representing not_indexed as proof that a person's data is **not on a site**, when it means only that it was **absent from the search index we read**.

If you resell Output, the three-state contract travels with it. The words the Service ships inside every response body — *"THREE STATES ONLY. This endpoint reads a search index; it does not fetch the page. not_indexed means absent from the search index — not proven absent from the site."* — are the words your customer is entitled to as well.

4.7 In violation of any applicable law

— or of a third party's terms of service, where you direct the Service at that third party.

Inside this prohibition, by example: naming, in a Verification, a domain whose terms forbid automated querying, where you are the party directing us there. The Service's own conduct toward third-party sites is described in Section 6 below; that conduct is ours and it does not make your instruction lawful.

In plain English. Seven rules. No judging anyone's eligibility for anything — that is background-check territory and we are not a background-check company. No hunting people. No faces or biometrics. No stolen material. No attacking or gaming the meter. No passing our answers off as a background check or a government record. Nothing illegal. Everything else is yours. And if you resell, your customers get the same seven, and so do theirs.

5. Obligations that sit alongside the fence

Nothing in this Section adds a restriction the Terms do not already impose. Each item names where it comes from, and is restated here because a compliance reader looking at one document should not have to guess. Breach of an item in this Section is handled as a breach of this Policy other than Section 4 — see Section 8 below.

- **Your representations about what you send** (Terms Section 7). For every Identifier you submit, you represent that you hold it lawfully, that you have a lawful basis to have us process it, and that you are not knowingly submitting the Identifier of a person you know or reasonably should know is under 18. **The Service does not screen for age and performs no date-of-birth boundary check. This representation is the control, and it is yours.** Nothing in this Policy or in the Terms obliges Sirveil to detect, verify, or refuse an Identifier relating to a minor, and nothing here should be read as describing a control that exists.
- **Key hygiene** (Terms Section 5). One key, one organization. Keys may not be shared across legal entities, sold, sublicensed, or embedded in a client application where an end user can extract them. An account may hold up to ten unrevoked keys. Tell us at support@sirveil.ai immediately if a key is exposed.
- **Use-case attestation** (Terms Section 5). We **may** require any account to certify its use case in writing, at key issue or afterwards, and may condition, suspend, or refuse access on that certification or on a refusal to give one.
- **Not for competing use** (Terms Sections 12(a) and 13). You will not use the Service, the Documentation, or the Output to build, train, or enrich a dataset, model, or service that detects, scores, or reports personal-data exposure in competition with the Service; to replicate the Registry; or to conduct benchmarking or competitive analysis for, or on behalf of, a competitor. Sell the answers; do not assemble them into the thing that answers.
- **Diagnostics are not determinations** (Terms Section 9.1). The confidence, calibration and linkage values in an Answer are returned to explain an answer, not to measure its accuracy. You will not build a decision threshold on one, and you will never put one in front of the person it is about.

6. What Sirveil does on its own side

This Section describes controls that exist in the code today, because you are entitled to know what the machine you are buying actually does. **It is a description of how the Service operates, not a warranty of how it will always operate.** Section 15 of the Terms states what Sirveil warrants, and Section 15.1 states what Sirveil does not have. Nothing in this Section creates a service level, and Section 11 of the Terms states that no uptime commitment is offered.

Facial recognition is refused at the boundary. A request naming a domain in the facial-recognition and biometric-identification class is refused with 403 `domain_excluded` **before any query is built and before any money is spent.** The refusal runs on **both** endpoints — the Verification path and the Scan registry — and it is layered, including a **build-time guard that fails the build** if the exclusion is removed. It is a control in the code, not a policy applied by hand, and **it is not waivable by anyone, at any price.**

The crawler tells the truth about itself. The username prober — the only part of the Service that requests pages from third-party sites directly — **fetches, parses and honors** `robots.txt`, **and fails closed.** It treats every ambiguous outcome as Disallow: a 401, a 403, a 5xx, a timeout, and a `robots.txt` it cannot parse. It sends an **honest identifying user agent carrying a contact URL** that explains what it is and on whose behalf it runs. It performs **zero retries**, and it **truncates a fetch at 32 KB.**

That describes **the Service.** It is not a representation about Sirveil's other internal systems, which are not part of the Service and are not sold to you.

There is no credential use anywhere in the product. No logging in, no captcha solving, no login bypass, no circumvention of a technical access control — none of it exists anywhere in the tree. The breach and stealer-log lane reports **whether** an identifier appears in a corpus; **nothing in the product uses a breached credential** for any purpose.

A Verification reads an index; it does not fetch the page. That is why `not_indexed` means absent from the search index and not proven absent from the site, and it is why the Service does not need to fetch broker pages directly. It does not.

You may not ask us to disable any of the above, and we will not.

***In plain English.** We will not check a face-search site for you, and that refusal happens in code before we spend a cent — nobody can buy their way past it. Our prober says who it is, leaves a contact URL, reads `robots.txt`, and treats "I could not read it" the same as "no." We never log in anywhere, never solve a captcha, never step around a lock, and never touch a stolen password. This paragraph describes what the code does today. It is not a promise about next year, and we are not going to dress a description up as a guarantee.*

7. If you resell, or make Output available to others

Section 7 of the Terms requires the fence to travel with the data. Concretely, you will:

1. **Bind your customers by written contract** to restrictions at least as protective as Section 4 of this Policy and Section 7 of the Terms;
2. **Require each of them to impose the same restrictions on every further recipient**, so that the fence survives more than one hop;
3. **Name Sirveil as an intended third-party beneficiary** of those restrictions, entitled to enforce them directly against your customer and against every further recipient. Section 24 of the Terms carries the matching provision on our side;
4. **Remain responsible to Sirveil for your customers' conduct**, as if it were your own; and
5. **On request, identify the categories of customer you serve** — categories, not names, unless a specific investigation requires more.

If a downstream customer breaches, **we may require you to cut them off as a condition of your own continued access.** We would rather you did that than we suspended you.

***In plain English.** The seven rules follow the data. Put them in your contract, make your customer put them in theirs, and say in writing that we can enforce them ourselves. If one of your customers turns out to be the problem, you cut them off — that is the version where you keep your account.*

8. How we enforce — the ladder

We have five rungs, and we generally climb them in this order:

#	Step	Where it comes from
1	Ask. We tell you what we saw and ask you to explain the pattern.	Terms Section 7
2	Require certification. We require a written use-case certification, and may condition access on it.	Terms Section 5
3	Throttle. We reduce a rate limit, a concurrency limit, or a quota. A refusal under a limit is not a breach by us and bills you nothing.	Terms Sections 11 and 11.1
4	Suspend. We suspend a key, an account, or a Marketplace entitlement immediately, without liability, on a reasonable belief of breach. A suspended account is refused 403 <code>tenant_suspended</code> .	Terms Section 18
5	Terminate. Immediately, without a cure period, for a breach of Section 4 of this Policy or Section 7 of the Terms. On 15 days' notice to cure for any other breach of this Policy.	Terms Section 19

We are not required to climb the ladder in order. Where the risk is immediate we may suspend or terminate first and tell you after. Skipping a rung once is not a waiver of it, and using a lower rung is not a representation that the conduct was acceptable.

What termination for a Section 4 breach costs you, stated plainly. Where we terminate for it, your license to Output ends as to further reproduction, distribution and resale; you will stop distributing Output and delete it on request; and you may keep using Output already embedded in a delivered product only so far as necessary to avoid harm to a third party, and not to make new copies (Terms Section 19). Separately, and whether or not we terminate, written notice of a Section 4 breach that you do not cure within 15 days ends your license as to Output delivered after that notice (Terms Section 12(b)). The liability cap does not protect you for that breach (Terms Section 17(c)). Fees already incurred remain payable. Where you purchased through a Marketplace, that Marketplace's published refund and cancellation policy applies to the transaction and nothing in this Policy narrows it.

We may decline to serve a use case we are not comfortable with, including a lawful one. That is not a determination that your use is unlawful, and we will say so if you ask.

We may monitor. We are not obliged to. We may monitor usage patterns for signs of prohibited use. **Nothing in this Policy or in the Terms obliges Sirveil to monitor, to detect misuse, or to prevent it, and no failure to detect misuse waives any right, excuses any breach, or constitutes approval of any use.** Serving a call is not a review of it. An account that has run for a year without a question has not been vetted.

***In plain English.** Most of the time this starts as an email asking what we are seeing. If the answer is good, that is the end of it. If it is not, we can throttle you, suspend you, or cut you off — and for the seven items in Section 4 we can cut you off the same day, keep the fees, and stop your right to resell what you already pulled. We are also not your compliance department: if we never noticed something, that is not us saying it was fine.*

***One channel.** Every message to Sirveil — privacy requests, security reports, abuse reports, legal notices — goes to support@sirveil.ai. Put the matter in the subject line (**Privacy · Security · Abuse Report · Legal Notice · Remedy Petition**) and it reaches the right person on the first read. Your message is effective whether or not the subject line is right.*

9. Reporting misuse

If you believe the Service is being misused — including against you — write to support@sirveil.ai. Those are read by a person. Include what you saw, when, and any request or response identifiers you have; every response, **including every error response**, carries the version stamps that produced it, and those make an incident traceable.

A **Subject** — a person whose information may have been queried — can reach us at support@sirveil.ai. Section 10 of the API Privacy Notice explains what we can and cannot do for a Subject, and it is honest about the limit: a synchronous call persists nothing, so there is generally no dossier for us to produce or delete.

Reporting misuse in good faith is not itself a breach of this Policy. That is not a license to go looking: Section 4.5 requires our prior written agreement before any security, load, or penetration testing, and reporting a finding afterwards does not supply that agreement in advance. Where you come across something without testing for it and tell us promptly and privately, we will weigh that — as a matter of judgment under Section 8, not as a right you can assert.

10. Document control

Field	Value
Version	2.0
Date	2026-08-25
Basis	Technical deep dive 2026-08-24 (code landing/v34 @ 5ca896e; live OpenAPI 1.4.0; production data plane queried 24 Aug 2026) · Alignment brief, API legal set v2 · API Terms of Use v2.0, 2026-08-25 · Input Requirements Brief 2026-08-19 · Status of Record 2026-08-20
Incorporated by	API Terms of Use, Section 3 (item 3 in the order of precedence)
Effective date	25 August 2026

Drafting principle applied throughout this version. Describe generously; commit narrowly. Section 6 states every protection the product actually delivers — the boundary refusal, the fail-closed robots.txt handling, the honest user agent, the absence of any credential use — because a buyer should be able to read this Policy and see that the values are in the code. Each is stated as a **description of how the Service operates**, not as a warranty of how it will always operate, and Section 6 says so in terms.

Brackets closed in this version.

v1 bracket	Resolution
Cure periods	Closed. Immediate termination for Section 4; 15 days' notice to cure for other breaches of this Policy. Matches Terms Section 19 exactly.
Notice mechanics	Closed. Terms Section 24 (Notices) governs; enforcement notice goes to the account email.
Refund treatment on termination for breach	Closed as to this Policy: fees already incurred remain payable (Terms Section 19), and a Marketplace's own published refund policy applies to the transaction and is not narrowed (Terms Section 3). The listing-conformity question is Terms open item 4 and is not re-opened here.
Channel list in Section 2	Closed. Section 2 no longer names Azure or an aggregator; it points to Terms Section 3.1, which names AWS Marketplace and says so is the whole list.

API PRIVACY NOTICE — SIRVEIL EXPOSURE API

*Version 2.0 · Effective 25 August 2026 · supersedes the versions of 2026-08-21 and 2026-08-19 and the page previously published at sirveil.ai/scan-api/privacy. **Basis:** the technical deep dive of 2026-08-24 (code at [landing/v34 @ 5ca896e](#), live OpenAPI 1.4.0, production data plane queried 24 Aug 2026), the Input Requirements Brief of 19 Aug 2026, and the Status of Record of 20 Aug 2026. Where this document and the code disagree, the code is the fact and this document is the defect. Reviewed and settled 25 August 2026. Every question carried by the previous version is decided and written out; nothing in this notice is provisional. Every blocking bracket carried by v1 is closed and written out. Incorporated by reference into the **API Terms of Use**, Section 3. Capitalized terms not defined here have the meaning given in Section 2 of those Terms.*

About the plain-English notes. Under most sections there is a short note explaining what that section means. The notes exist to help; they are not legal advice, they do not replace the formal text, and where a note and the legal text differ, **the legal text controls**.

1. The short version

You send us identifiers about a person. We look them up in sources we can reach. We return an answer. Most of what follows is detail on that sentence, but four things are worth knowing before the detail, and one of them is uncomfortable.

First: which endpoint you called changes where the data goes, and we are not going to blur the two together.

- **A Verification reaches exactly one third party — the search provider — and makes zero model calls.** No enrichment vendor. No breach vendor. No court-records source. No username site. One `site:-`scoped query, one recipient, one answer.
- **A Scan reaches more.** In addition to the search provider it may reach identity enrichment, breach and stealer-log lookup, court records, and public username pages, and it uses model inference on tokenized data for parts of its adjudication and phrasing.

Second: some of what a Scan reaches is data we pay for. The identity-enrichment vendor and the breach-corpus vendors are **commercial subscription data**, not public web pages. That is true, it is in Section 5, and it is stated there rather than softened.

Third: what we keep, in three cases. A synchronous Verification or Scan **does not persist the identity you sent or the answer we returned**. An asynchronous job **holds both for 24 hours**, then the sensitive columns are emptied. And every authenticated call writes **one activity row** — who called, which endpoint, how long it took, what it billed — **containing no request body and no response body**, and therefore nothing about the Subject. That row has **no expiry today**. We would rather write that sentence than tell you we keep nothing.

Fourth: two commitments that are permanent. We do not sell the Identifiers you send us or the Output. We do not use them to build, seed, train, or enrich a separate product or a people database. Those two are floors, not snapshots, and a later version of this notice will not weaken them retroactively for data already collected.

2. What this notice covers

This notice covers the **Sirveil Exposure API** — every call to `POST /api/v1/verify`, `POST /api/v1/scan`, `GET /api/v1/jobs/{id}`, `GET /api/v1/whoami`, and the read-only agent surface at `POST /api/v1/mcp` — and the account and billing data that goes with it.

It does **not** cover the Sirveil consumer application, which has its own separate Privacy Policy. That policy is not modified by this one, and this one does not govern the app.

Where this notice and the Terms differ on a subject the Terms address, **the Terms control** (Terms Section 3). Where this notice and any Sirveil marketing surface, listing, or datasheet differ on what the Service does with data, **this notice controls**. Section 5 says so specifically, and says why.

3. Roles

Party	Role
You, the API customer	Controller of the Identifiers you submit. You decide why they are submitted, you hold the lawful basis, and you owe the Subject the explanations.
Sirveil, as to those Identifiers	Service provider (as defined by the CCPA) and processor . We use them to answer your call and for the limited retained purposes in Section 6, and for no other purpose.
Sirveil, as to your account	Independent business with respect to your account data, billing data, and abuse-prevention data <i>about you</i> — not about your Subjects.
Our suppliers	Subprocessors , engaged by category as described in Section 8, bound to use the data only to provide their service to us.

The service-provider terms are in the agreement itself. Section 8 of the API Terms of Use carries the covenants the California Consumer Privacy Act requires of a service provider — no sale, no sharing, no use outside the stated business purposes, no combining with other sources, notice to you if we can no longer meet those obligations, and your right to verify and to stop and remediate. **No separate data processing addendum is needed for Sirveil to act as your service provider.** Where your own regulator or customer requires an addendum on their paper, write to support@sirveil.ai and we will review it.

***In plain English.** You decide why a lookup happens; we run it for you and nothing else. The person whose data it is has a relationship with you, not with us — which is why Section 10 says to send them to you first. Your own account and invoice data is a different thing, and there we are our own business, not your processor.*

4. What we collect

From a call — about the Subject

Required, on both endpoints: firstName and lastName. A Verification additionally requires a domain. The name is the anchor every candidate record is scored against; without it the Service has no way to tell whether the record it found is the right human being. A request missing a required field is refused with 400 identity_incomplete **before any lookup runs and before anything is charged.**

Optional, and entirely your choice: phone number, email address, street address, city, state, date of birth, and one or more usernames. The more of these you send, the more the query planner has to work with — it leads with a held identifier where it has one, in the order **phone** → **email** → **street address** → **name**. Supplying a username is what switches on the username lane at all; the Service does not guess a handle from an email address.

A field we cannot parse is dropped rather than guessed at, and the name of the dropped field is returned to you in dropped_fields.

What you must not send. Passwords, password hashes, session tokens, full payment-card data, government-ID images, face images, face embeddings, biometric templates, and data from a breach or dump you have no lawful right to hold. That is Terms Section 7(4) and Section 4.4 of the AUP, and it is a contractual prohibition, not a filter — **the Service does**

not inspect your payload for prohibited content and will not catch it for you.

From every call — for metering and abuse prevention

Every authenticated call writes exactly **one activity row**, from the Service's own accounting path. It contains:

- the account and the key that made the call (the key by its stored identifier, never its secret);
- the endpoint, the timestamp, the HTTP status, and the latency;
- the billable units and the outcome (served, unserved, refused);
- an internal run identifier, which is what makes an answer traceable to the logic that produced it; and
- an internal cost figure of Sirveil's own — what the attempt cost **us**, not a price to you, and not information about the Subject.

It contains no request body and no response body. No Identifier and no Answer is written to it. That row is the billing record, and an invoice you can audit depends on it existing.

Refusals are accounted for too, and honestly. An **authenticated** call that is refused writes a row with an outcome of refused, a refusal code, **zero** billable units, and a **null** cost — deliberately null rather than zero, because zero would assert a measurement we never made. An **unauthenticated** refusal — a missing or invalid key — writes **no row at all**, because the table is keyed on an account; that volume exists only in ordinary request logs.

About your account

A business email address, an organization name, the records of the keys we issued (a hash, a short non-usable prefix, and a revocation state), and the invoice ledger.

Billing is handled by your purchase channel. Payment is taken by that channel and **we never see or store card numbers.**

In plain English. We take what the lookup needs and nothing else. Every call also leaves one line in our ledger — who called, which endpoint, how long, what it cost us, what it billed you — and that line has nothing about the person in it. It is how your invoice is auditable, and it is the only thing that survives a synchronous call.

5. Where the data goes — and Verification and Scan are not the same

This is the section a security reviewer should read first, and it is the section most privacy notices for a service like this get wrong by averaging two different products into one sentence. We are not going to do that.

Category of recipient	Used by	What it receives
Search / retrieval	Verification and Scan	The site-scoped query, containing the Identifier in clear text
Identity enrichment	Scan only	Identity fields in clear text. Commercial subscription data.
Breach and stealer-log lookup	Scan only	The Identifier in clear text; for some sources an email address appears in a request URL path. Commercial subscription data.
Court records	Scan only	A name-based query
Public username pages	Scan only	Only a handle you supplied, in a profile URL

Category of recipient	Used by	What it receives
Model inference	Scan only	Tokenized placeholders only — never a raw name, phone number, email address, street address, date of birth, or government identifier

A Verification reaches exactly one third party and makes zero model calls. One `site:-`scoped query to the search provider, and it stops there. A failover search provider may stand in where the primary fails; that is the same single category of recipient and the same single call, not an additional one. Nothing about a Verification reaches an enrichment vendor, a breach vendor, a court-records source, a username site, or a model.

A Scan reaches more, and some of what it reaches is data we buy. The identity-enrichment vendor is a purchased-data enrichment service. The breach and stealer-log corpora are commercial subscription services. These are paid vendors, not public web pages, and a Scan queries them on your behalf.

This Section controls over any contrary statement on a Sirveil surface. If a Sirveil marketing page, listing, or datasheet says or implies that answers come from publicly indexed pages only, with no purchased data about your subjects, **that statement is accurate for a Verification and is not accurate for a Scan**, and this Section is the one that describes the Service. The correction to those surfaces is tracked in Section 15 as a conforming change.

****Breach-lane vendors report *whether*, not *what*.** The breach and stealer-log lane returns **whether an identifier appears in a corpus**. Nothing in the product uses, returns, replays, or tests a breached credential**, and nothing in the product logs in anywhere, solves a captcha, or circumvents an access control. AUP Section 6 describes that conduct in full.

The username lane behaves. It requests only public profile URLs for a handle you supplied. It fetches, parses and honors `robots.txt`, **failing closed** — treating a 401, a 403, a 5xx, a timeout, and an unparseable file as Disallow. It sends an honest identifying user agent with a contact URL, performs zero retries, and truncates at 32 KB.

In plain English. A check goes to one search provider and stops. A sweep goes to more places, and some of them are vendors we pay for data — an enrichment service and breach corpora. If you read somewhere on our site that we never use purchased data about your subjects, that sentence is right about a check and wrong about a sweep, and this page is the one that is right. The breach lane only ever tells us whether an address shows up in a dump. It never touches the password.

5.1 The PII firewall — what it protects, and the gap it does not close

Before any model call, sensitive values are replaced with placeholders: a name becomes `PERSON_1`, a phone number becomes `PHONE_2`, an address becomes `ADDR_1`. The protected class is **name, phone number, email address, street address, date of birth, government identifier, and exact geolocation** — none of these may reach the model layer.

The control is not a convention. A check runs on the outbound payload, and **if a raw protected value has survived tokenization the check throws and the outbound call is aborted rather than sent**, with a per-call trace recorded. A model call that would have carried a real identifier does not happen; it fails.

Model routing is **pinned to a US-host allowlist**, configured with **data collection denied** and **zero data retention requested** on every provider. Providers hosted outside the United States are excluded by policy. Section 12 carries the open question about the standing of that policy.

The gap, stated because you would find it anyway. Tokenization protects the model path only. The retrieval vendors — search, identity enrichment, breach and stealer-log lookup, court records — **necessarily receive the Identifier in clear text, because that is what a lookup is**. There is no way to ask an index whether a phone number appears on a page without sending the phone number. We are not going to pretend otherwise, and a vendor who tells you their retrieval

works differently is describing a different operation.

This describes how the Service operates; it is not a warranty. Terms Section 15 states what Sirveil warrants and Section 15.1 states what Sirveil does not have. And note the asymmetry that follows from Section 5: because a Verification makes **no model calls at all**, this firewall is not exercised on a Verification — there is no model path on that endpoint to protect.

In plain English. The AI models only ever see PERSON_1 and PHONE_2. If a real value ever leaked into that path, the call gets killed instead of sent — that is code, not a policy. But be clear about what that does and does not buy you: the lookup vendors get the real phone number, because a lookup with a fake phone number is not a lookup. Tokenization protects the model side. It does not protect the retrieval side, and nothing can.

6. What we do with it

We use the Identifiers you submit **solely** to run your query against the sources the Service is built to query, to shape the answer, and to return it to you.

We use activity and account data to bill accurately, to enforce the limits described in Terms Section 11.1, to prevent and investigate abuse under the AUP, to support you, and to measure our own performance.

We do not:

- sell the Identifiers you submit or the Output;
- share them for cross-context behavioral advertising;
- use them to build, seed, train, or enrich a separate product or a people database;
- enrich or append the identities you send us for our own purposes;
- combine them with data from another source except as needed to answer your call; or
- use your query patterns to compete with you or to inform anyone else's product.

The first three are permanent as to data already collected, and a later change to this notice will not weaken them retroactively. As to the rest, we will not weaken them for data already collected without your consent. None of this stops us using aggregated, de-identified signals about abuse, security, and performance to run and protect the Service.

A cached answer is still processing, and it is still billed. A short-lived query cache holds normalized query text — which contains the Identifier — in memory only, partitioned by account, for minutes rather than hours. Serving from it is still processing of the Identifier, and Terms Section 10 says plainly that a cache hit bills and consumes quota exactly like a fresh call. Neither document buries that.

7. Retention — all three cases, stated exactly

Store	What it holds	How long
Synchronous Verification or Scan	Nothing about the Subject is written. The Verification path is guarded in code against database writes.	No Subject data is persisted
Asynchronous job (Prefer : respond-async only)	The submitted identity and the shaped result — it has to, between submit and poll	24 hours. Then the sensitive columns are nulled and only the job's id, status and timings survive, so a late poll receives an explicit expiry rather than an answer indistinguishable from a bad id

Store	What it holds	How long
Per-call activity row	Account, key identifier, endpoint, timestamp, status, latency, billable units, outcome, run identifier, internal cost figure. No request body. No response body.	Indefinite today — see below
Query cache	Normalized query text, containing the Identifier; per-instance, in memory only, partitioned by account	Minutes (currently 15), with a bounded entry count. Not written to disk
Account and billing records	Business email, organization, key records (hash, prefix, revocation state), invoice ledger	For the life of the account, and afterwards for as long as tax, accounting, and dispute-resolution law requires
<code>robots.txt</code> cache	A per-host parse result	In process. Not personal data

The activity row, stated honestly. There is **no expiry column on that table and no purge routine**. That is **an unresolved decision, not a designed choice** — nobody chose indefinite retention for that row; indefinite is what exists because no period was ever set. A stated period would be better disclosure than this paragraph, and setting one is on the list. Until it is set, this is what is true, and we would rather say so than write that the API accumulates nothing.

We do not keep asynchronous results as a convenience archive. When the 24 hours are up, the identity and the result are gone and there is no restore. A poll after expiry returns an explicit `result_expired`, not a silence you would have to interpret.

Revocation does not delete a key row. Revoking a key sets a revocation timestamp; the row stays, so historical usage can always name the key that produced it. That is deliberate — it is what makes an invoice auditable — and it is described in Terms Section 5.

***In plain English.** Call us and wait for the answer, and we do not keep the person's details or the result. Use the `async` option and we have to hold both until you come back — 24 hours, then they are wiped and only the timing shell of the job is left. Separately, every call leaves one line in our billing ledger with nothing about the person in it, and that line has no delete date yet. Nobody decided that. It is just what is true, and telling you is better than a sentence about keeping nothing.*

8. Subprocessors

We disclose subprocessors **by category rather than by name**, because a named list is also a map of our supply chain and it goes stale the moment a vendor changes.

The categories are: search and retrieval · identity enrichment · breach and stealer-log lookup · court records · model inference · username probing · cloud infrastructure and database hosting · billing and payment channels.

Section 5 states what each category receives. Each is engaged under contract to use the data only to provide its service to us.

The current list of named subprocessors is available on request at support@sirveil.ai, and we will provide it to a Marketplace reviewer or an enterprise buyer under a confidentiality undertaking. We will give notice of a material change to the categories. Sirveil may add, remove, or replace a subprocessor within these categories.

Two shapes worth naming, without naming the vendor. For some breach-corpus sources, **an email address appears in a request URL path** rather than as a hashed or k-anonymized lookup — that is the vendor's interface, not our choice, and it is disclosed here because a security reviewer will want to know. And **model-inference vendors receive tokenized placeholders only**, behind the fail-closed control in Section 5.1.

Vendor names are deliberately not printed here — a published list is also a map of our supply chain. They are available to you on request at support@sirveil.ai.

In plain English. We tell you the kinds of vendors, not the brand names, and you can have the brand names by asking. One thing we will say out loud: for some breach lookups the email address travels in the web address itself, which is how that vendor's interface works. It is the kind of detail we would rather you heard from us than found.

9. Security

Keys. An API key is generated from **192 bits of randomness** from a cryptographically secure generator. It is stored as a **SHA-256 hash only**, together with a short non-usable prefix for identification. The plaintext is **displayed exactly once, at issue**, and is never stored — we cannot recover it, re-display it, or email it to you. An account may hold a maximum of **ten unrevoked keys**. **Revocation never deletes the key row.**

Database access. Row-level security is **enabled on the account, key, activity and job tables with an empty policy set** — the safest policy set for a billing table is the empty one — and access is **service-role only**. Supporting operational tables additionally carry explicit revokes from the anonymous and authenticated database roles.

Responses are built by whitelist. Both public response shapers build every field **explicitly and never spread an internal object**, so vendor names, internal cost figures, model identifiers, build identifiers, raw scoring inputs, and provider-configuration flags cannot reach a customer payload by being added upstream. An independent gate asserts that no vendor name, cost field, or build identifier appears in a Scan response.

Cross-tenant correlation resistance. The internal correlation label — a stable hash over a Subject's identifiers, which would otherwise be identical across two accounts querying the same person — is **dropped** before a response is built. Scan identifiers and card identifiers are **minted per response**. Match probabilities are **rounded**, so internal cluster boundaries cannot be recovered from a distribution of answers. **Two customers querying the same person cannot learn that fact from anything the Service returns.**

Tenant isolation. A job identifier belonging to another account returns **404, never 403** — we do not confirm the existence of someone else's job.

Encryption, stated exactly. Data is **encrypted in transit with TLS**. Encryption at rest is provided by our hosting and database platforms and **has not been independently confirmed in either vendor console**. **There is no application-level field encryption**, and nothing in the product claims otherwise.

No certifications, no audit, no penetration test. Sirveil holds **no SOC 2 report, no ISO 27001 certification, and has not had a third-party penetration test**, and holds no other security or privacy certification, attestation, or audit report for the Service. Terms Section 15.1 states the same thing and states it as a term. If your organization requires one as a condition of purchase, the answer today is that we do not have one — ask before you subscribe.

Incidents. We will notify you **without undue delay** after we confirm a security incident affecting data you submitted or your account, and will give you the information you reasonably need for your own notification duties. **We do not state a fixed hour count here.** A number in a published notice is a commitment measured against operations, and we would rather give you a real one in a negotiated agreement than a round one on a web page. Ask, and it is negotiable.

In plain English. Your key is a password we genuinely cannot get back for you — a key we could recover is a key someone could talk us out of. The billing tables are locked to the service role with no policy that opens them. Responses are assembled field by field, so our vendors and our costs cannot leak into your payload by accident. Two of you scanning the same person cannot tell. And the honest half: TLS yes, at-rest encryption from our platforms that we have not gone and verified in the console, no field-level encryption, no SOC 2, no ISO, no pen test. If you need those, we are not ready for you yet.

10. Rights — yours, and your Subjects'

Yours, as the customer

Ask us for a copy of what we hold about your account, correct it, or delete it, at support@sirveil.ai. We respond within **30 days** of a verified request, at no charge.

Two honest limits. The per-call activity row described in Section 4 is the **billing record**; we retain it, and Terms Sections 8 and 19 say so. Key rows are not deleted on revocation, by design. Neither contains an Identifier or an Answer.

Your Subjects'

****A person whose Identifiers *you* submitted deals with *you* first.**** You are the controller, you know why the query was run, and you hold the relationship. That is not a deflection — it is the allocation in Section 3, and it is the one the law contemplates.

The honest limit, said plainly: Sirveil generally holds no dossier to disclose or delete for a Subject. A synchronous call persists neither the identity submitted nor the answer returned. The activity row contains no Identifier, so it **cannot be searched by Subject at all** — there is no per-Subject store to look in. A request to "delete everything you hold about me" usually has nothing to act on, and we will say that rather than perform a deletion ceremony over an empty set.

What we can actually do for a Subject:

- tell them we act for a business customer, and explain this role honestly;
- refer them to that customer where we lawfully can;
- assist that customer in responding, on their instruction;
- where an **asynchronous job** containing their identifiers is still inside its 24-hour window, **purge it early on our customer's instruction** — this is the one case where there is something to delete. We will not confirm to a third party that a particular person has been the subject of a query; that is our customer's information, not ours to disclose; and
- point them to their own rights **against the source** — the site or broker that published the data, which is where the data actually lives — and to the **Sirveil consumer application**, which is built for a person checking their own exposure and is priced for it.

As a service provider we do not decide what our customer keeps, and we do not act on a request about a customer's processing without that customer.

A Subject can reach us at support@sirveil.ai. We read those.

What we do with a request that reaches us directly. We will (a) tell the Subject, promptly, that Sirveil processed the request on a customer's instruction and that the customer holds the record; (b) **forward the request to that customer** where we can identify them, rather than merely referring the Subject onward; and (c) where we genuinely hold nothing responsive, say so plainly rather than performing a deletion of something that does not exist. We will not knowingly leave a Subject with nowhere to go.

***In plain English.** If someone asks us to hand over or delete their file, the truthful answer is usually that there is no file — a synchronous call leaves nothing behind, and our billing ledger has no names in it to search. We will not pretend to delete something that does not exist. What we will do is tell them who we ran it for, help that customer answer them, wipe an in-flight async job early if asked, and point them at the site that actually published their data — and at our consumer app, which is the product built for exactly this.*

11. Children

The Service is not for looking up minors. You may not submit the Identifiers of a person you know or reasonably should know is under 18.

The Service does not screen for age and performs no date-of-birth boundary check; this representation is the control, and it is yours. Nothing in this notice or in the Terms obliges Sirveil to detect, verify, or refuse an Identifier relating to a minor, and nothing here should be read as describing a control that exists. Terms Section 7 states this in the same words.

12. Territory and data residency

The Service is offered **in the United States only**, and data is processed in the United States. **Model routing is pinned to a US-host allowlist** by configuration, with providers hosted outside the United States excluded by policy.

We do not currently offer the API in the EU or UK and make **no cross-border transfer commitments** for it. If that changes, a data processing addendum with transfer terms comes first — not after.

A Marketplace listing is visible worldwide even though this offer is not. The Service is offered to, and the Terms are written for, customers established in the United States, and it is processed in the United States. If you are established outside the United States, the Service is not offered to you and you should not subscribe; where Sirveil identifies such a subscription it may decline to provision and will refund it through the Marketplace's own process.

13. Not a consumer reporting agency

Sirveil is **not** a consumer reporting agency and the Output is **not** a consumer report. An Answer is a point-in-time observation of an index; it is not a statement about a person and it is never a determination about anyone's character, conduct, or eligibility.

The prohibited uses in Terms Section 7 and Section 4 of the AUP are contractual, and they are enforced under the ladder in Section 8 of the AUP.

14. Changes, and how to reach us

We will give at least **30 days' notice** to your account email before a material change to this notice. Changes apply prospectively. The commitments in Section 6 that are marked permanent — no sale, no separate people database — are floors, not snapshots, and a later version will not weaken them retroactively for data already collected. Each version carries a version number and an effective date.

Contact. Privacy: support@sirveil.ai · Security and incidents: support@sirveil.ai · Abuse and misuse: support@sirveil.ai · Legal, including subprocessor list requests and DPA inquiries: support@sirveil.ai · Support: support@sirveil.ai.

Sirveil, Inc., a Delaware corporation · San Jose, California.

Disputes follow **Terms Section 22** — a written remedy petition, 60 days to resolve it between people who can settle, then arbitration.

***One channel.** Every message to Sirveil — privacy requests, security reports, abuse reports, legal notices — goes to support@sirveil.ai. Put the matter in the subject line (**Privacy · Security · Abuse Report · Legal Notice · Remedy Petition**) and it reaches the right person on the first read. Your message is effective whether or not the subject line is right.*

15. Document control

Field	Value
Version	2.0
Date	2026-08-25
Supersedes	The versions of 2026-08-21 and 2026-08-19, and the page previously published at sirveil.ai/scan-api/privacy
Basis	Technical deep dive 2026-08-24 (code landing/v34 @ 5ca896e ; live OpenAPI 1.4.0; production data plane queried 24 Aug 2026) · Alignment brief, API legal set v2 · API Terms of Use v2.0, 2026-08-25 · Input Requirements Brief 2026-08-19 · Status of Record 2026-08-20
Incorporated by	API Terms of Use, Section 3 (item 4 in the order of precedence)
Effective date	25 August 2026

Drafting principle applied throughout this version. Describe generously; commit narrowly. Every protection the product actually delivers — the fail-closed PII firewall, the whitelist projection, the correlation-label drop, the empty-policy row-level security, the never-recoverable key — is **stated**, because a security reviewer should be able to read this page and see the architecture. Each is stated as a **description of how the Service operates**, not as a warranty of how it will always operate. Terms Section 15 carries the warranties; Section 15.1 carries what Sirveil does not have.

Brackets closed in this version.

v1 bracket	Resolution
Purchased data — "is purchased data on the API path?" (v1 § 6, blocking)	Closed, and the answer is the uncomfortable one. Yes, on a Scan. Section 5 describes it, separates Verification from Scan, and states that this notice controls over any contrary marketing sentence.
Async retention — 24 hours or 90 days (v1 § 5, blocking)	Closed at 24 hours , matching the code and Terms Section 8. The live page's 90-day statement is wrong and must come down.
Accepted identifier set (v1 § 3)	Closed: <code>firstName</code> and <code>lastName</code> required on both endpoints, <code>domain</code> additionally on a Verification; phone, email, street address, city, state, date of birth and usernames optional. The live page's "name, email, state" is narrower than what the API accepts and must be corrected.
Explicit revokes on the billing tables (v1 § 7, drafting error)	Corrected. v1 implied explicit <code>revoke</code> grants on the core account, key, activity and job tables. Those tables rely on row-level security with an empty policy set , service-role only; the explicit revokes sit on the supporting operational tables. Functionally equivalent, but the v1 sentence was wrong. Section 9 states it correctly.
Metering row contents (v1 § 3)	Closed and expanded: the full field list is in Section 4, including the internal cost figure and the run identifier, with the "no request or response body" point stated twice because it is the one that matters.
Certifications	Closed: Section 9 states no SOC 2, no ISO 27001, no penetration test, no application-level field encryption, and at-rest encryption not independently confirmed.

One decision is a business one, not a legal one. Section 7 discloses that the per-call activity row is retained indefinitely, because that is the current state. It is an unresolved decision rather than a designed choice, and a stated retention period would be better disclosure than none. That is a change to make in the product, and this notice will follow it.