COMPLIANCE BRIEF

What decides whether your AI agent can *legally call* someone.

The TCPA rules that govern automated calls and texts in the United States, mapped onto the two things that actually change your exposure: **who starts the conversation**, and **what your opt-in actually says**.

PUBLISHED BY	Hisora · hisora.ai
EDITION	One — September 2026
SCOPE	US federal law, plus the state rules that stack on top
STATUS	Working reference — not legal advice

FIRST PRINCIPLES

A 1991 privacy statute that never heard of AI — and governs it anyway.

Thirty seconds of background, because every rule that follows is a consequence of how the law was built — not of anything anyone decided about AI.

Congress passed the **Telephone Consumer Protection Act in 1991**, aimed at telemarketers, autodialers, prerecorded messages and junk faxes. It is codified at **47 U.S.C. § 227**; the FCC writes the implementing rules, at **47 CFR 64.1200**. Nothing in it was designed for conversational software — which is the problem. The definitions are broad, they are thirty-five years old, and an AI voice landed inside one of them.

What it restricts — keyed to what you dialled, never to what you meant

PROVISION	WHAT IT RESTRICTS	WHOSE NUMBER
§ 227(b)(1)(A)(iii)	Autodialed or artificial-voice calls and texts, without prior express consent	Any **cellular** number
§ 227(b)(1)(B)	Artificial or prerecorded voice calls, without prior express consent	**Residential** lines
§ 227(c)	Telemarketing to a number listed on the national Do Not Call registry	Registered **residential** subscribers
64.1200(b), (d)	Duties on the call itself — identify yourself, give a callback number, keep an internal do-not-call list	Everyone, **B2B included**

Two regulators, not one: the **FCC** enforces the TCPA, while the **FTC** runs the Do Not Call registry under its own Telemarketing Sales Rule — the source of the 31-day re-scrub duty.

Texts count, and the plaintiff is the person who answered

SMS did not exist in 1991 and is covered regardless: the FCC and the courts read a text as a "call" under § 227(b). A 2026 decision narrowed the *registry* route for texts in three states and left the artificial-voice route — the one an AI agent sits in — intact everywhere.

Enforcement does not wait for a regulator to notice you. The statute carries a private right of action with fixed per-call damages and a four-year window running from each individual violation, so the person who picks up the phone is the person who sues you.

WHERE THE TCPA STOPS

It keys on **calls to a telephone number**. WhatsApp and in-app chat route to an account rather than a number, so they sit outside it — which is not permission to blast them. **Meta's own policy is stricter than the TCPA in practice** (opt-in, pre-approved templates, a 24-hour service window), and state unfair-practice law does not care what the transport was.

START HERE

An AI voice is an artificial voice. Everything else follows from that.

Conversational quality is legally irrelevant. However human your agent sounds, the rulebook it lands in was written in 1991.

In **February 2024**, FCC Declaratory Ruling **24-17** confirmed that an AI-generated voice is an **“artificial or prerecorded voice”** under the Telephone Consumer Protection Act, effective immediately. There is no “but it’s an AI, not a robocall” argument.

$500
PER CALL OR TEXT, MINIMUM

$1,500
PER CALL, IF THE VIOLATION IS WILLFUL

5 yrs
INTERNAL DO-NOT-CALL RETENTION REQUIRED

The regulator is not the threat

The TCPA carries a **private right of action**. Anyone who received the call can sue, and their attorney can certify a class straight out of your own dial log — the record you kept to prove compliance is the same record that sizes the claim.

THIS STOPPED BEING HYPOTHETICAL IN 2026

Two proposed class actions were filed fifteen days apart in the Eastern District of Michigan over AI voice cold calls — *Lamb v. Mortgage One Funding* (2:26-cv-10450) and *Landy* (2:26-cv-10643). **The proposed classes name the company’s vendors, lead generators and agents** — not only the business whose name was on the call. Both complaints also put *“it never disclosed it was AI”* directly beside the consent claim.

Why the ground is still moving

Three decisions in eighteen months make anything written before 2025 stale, and they point one way: **courts are stripping the FCC’s add-ons while the statute, where the damages live, has not moved.**

Insurance Marketing Coalition v. FCC 11th Cir. · Jan 2025	Vacated the one-to-one consent rule
McLaughlin v. McKesson SCOTUS · Jun 2025	District courts are no longer bound by the FCC’s reading of the statute
Bradford v. Sovereign Pest Control 5th Cir. · Feb 2026	Oral consent valid in TX, LA, MS — three states out of fifty

Written consent remains the only posture that works in all fifty states.

THE DISTINCTION THAT DECIDES EVERYTHING

Who dialled first is the single biggest variable in your exposure.

Two systems can use identical technology and sit in completely different risk positions. The difference is which end started the conversation.

PATTERN	WHAT IT IS	EXPOSURE
Inbound	A customer calls you. The agent answers, books, answers questions.	**Minimal. They initiated it.**
Outbound — transactional	Appointment confirmations and reminders to someone who just booked. Consent arises from the transaction itself.	**Low, and conditional**
Outbound — warm lead	Someone filled in your form asking to be contacted, and the agent calls them back.	**Defensible — if the opt-in named the right technology**
Outbound — marketing	Review requests, reactivation, promotions, anything with an offer in it.	**Prior express *written* consent**
Outbound — cold	A list they never asked to be on.	**The lane the 2026 lawsuits are about**

THE RULE THAT CATCHES CAREFUL PEOPLE

One upsell reclassifies the entire call. The FCC treats a dual-purpose call as marketing. So an otherwise compliant appointment reminder that adds *"and while I have you, we're running 20% off this month"* becomes a telemarketing call requiring written consent nobody collected. On our builds this is a **hard prohibition in the agent prompt and a test case**, not a guideline — because the sentence that breaks it is exactly the sentence a good salesperson would add.

And the number decides, not the person

There is a widely repeated belief that calling businesses is safe. It is half right, and the half that is wrong is expensive.

The Do Not Call registry and the residential-line restriction both key off *residential*, so business landlines genuinely are an open lane. But **47 U.S.C. § 227(b)(1)(A)(iii) covers any number assigned to cellular service, full stop** — the statute never asks whether the person holding it is a consumer or a business owner.

For contractors, clinics and owner-operated businesses, **the number on the website is usually the owner's mobile**, and **76% of US adults have no landline at all** (CDC). After number portability, area code and prefix prove nothing. That is what turns a "B2B only" campaign into a class action.

THE OPT-IN

A checkbox is a signature. The words around it are the job.

47 CFR 64.1200(f)(9) defines prior express written consent. Miss an element and what you are holding is not consent — it is a screenshot.

The six elements

> "I agree that **[COMPANY NAME]** may contact me at **the phone number I provided above** using **an automatic telephone dialing system or an artificial or prerecorded voice, and by text message**, for marketing purposes. **I understand consent is not a condition of purchase. I can revoke consent at any time** by replying STOP or telling the caller to stop."

The named company · the specific number · **the technology** · the purpose · not-a-condition-of-purchase · the revocation route.

THE ELEMENT EVERYBODY SKIPS

The technology clause decides whether your record covers an AI agent at all. It has to say autodialer *and* artificial or prerecorded voice. A form that says only "calls and texts" does not cover it. **Text consent is not voice consent** — they are separate authorizations, so a clean SMS opt-in list is telling you the truth about texts and saying nothing at all about calls.

Proving it is your job, not theirs

Consent is an **affirmative defense**. The plaintiff never has to prove they did not consent; you have to prove they did. Until you produce a real record you are losing the argument. **Which makes the record the product.** Five fields per lead, stored somewhere queryable:

1. **Timestamp** of the submission
2. **IP address** it came from — not required by rule, and the only answer to *"somebody else submitted my number"*, which is the most common way a consent record dies
3. **Checkbox state**, captured rather than assumed
4. **The phone number exactly as entered**
5. **The version ID of the disclosure they saw** — you must prove what was on screen that day, not what the form says today. Version the text and store the ID with the record.

Consent you did not collect yourself

Purchased lists, lead aggregators, affiliate traffic, a ten-year-old CRM export. Same question every time: did the third party capture written consent naming *this* company and *this* technology? If nobody knows, it does not matter who collected it — **you placed the call, so you are the defendant.** A TrustedForm or Jornaya certificate is not consent, it is *proof* of consent: pull one and read what is inside it.

DUTIES AND REVOCATION

Three duties ride along on every artificial-voice call — even B2B.

Identify	The agent states the business name **at the start**. Your opening line is a compliance surface, not just a greeting.
Callback	A real number the recipient can call back, given on request.
Internal DNC	Your own suppression list, retained **five years**, fed by *every* channel — including the support inbox and anything said out loud on a call.

Revocation works by any reasonable means

You cannot designate an exclusive opt-out method. **"Take me off your list" spoken to the voice agent is a valid revocation**, and no keyword matcher will ever catch it. Removals must propagate within **10 business days**, in force since April 2025.

The FCC's broader rule — one revocation reaching every system you own that can dial or text — has a compliance date of **31 January 2027**. We build to it now, because a cross-channel suppression table is not something anyone wants to retrofit under deadline.

HOW WE HANDLE IT IN PRACTICE

Keyword matching (`STOP`, `UNSUBSCRIBE`, `CANCEL`) catches the easy cases, and it is **not sufficient on its own**. Two things cover the rest: the agent is prompted to recognise and honour a spoken opt-out live, and **any inbound reply the matcher does not recognise is routed to a human rather than dropped.** That routing is a compliance control, not a convenience.

The mirror-image error costs money instead of damages: *"can you call me back next week"* is a reschedule, not an opt-out. Suppressing it loses a live lead.

Disclosure, recording and quiet hours

- **There is no federal AI-disclosure rule.** California's **AB 2905** (effective 2025) requires it up front in a natural voice; Utah requires it on request. Both 2026 complaints raised it anyway, so we disclose by default.
- **All-party consent states matter here** because voice platforms typically record every call by default. That is a setting to make deliberately, not to inherit.
- **Calling and texting hours are 8am–9pm in the recipient's timezone**, not the server's and not the business's. A two-hour reminder for an 8am appointment is a 6am text to someone one timezone west. Both timezones are needed at once: quiet hours where the phone is, appointment times where the business is.

BEFORE THE AGENT IS SWITCHED ON

Ten checks. If you cannot tick all ten, the campaign does not launch.

One to three is *who you may call*. Four to seven is *permission*. Eight to ten is *the call itself*.

- ☐ I know which numbers are consumer and which are business
- ☐ I know which of them are **mobile**
- ☐ Consumer numbers are scrubbed against the national DNC list, plus any state list
- ☐ Written consent is on file for every number that needs it
- ☐ The consent names **my company, that number, and artificial voice**
- ☐ I can pull the record for any single lead, **including the language they saw**
- ☐ If someone else built the list, their consent evidence is in my contract
- ☐ The agent gives the company name, a callback number, and discloses recording where required
- ☐ The agent honours *"take me off your list"*, and removals reach every system within 10 business days
- ☐ The dialer enforces **8am–9pm where the phone is**

What this looks like in the build

None of the above is paperwork that sits beside the system — each line is a component:

Consent table	Five fields per lead, versioned disclosure text, queryable by number
Suppression table	Cross-channel, checked before *every* send with no exception path. An SMS opt-out stops calls too
Quiet-hours gate	Computed in the recipient's timezone, applied inside the single send function so nothing can route around it
Message ledger	Deterministic key per lead, touch and channel — makes a duplicate contact structurally impossible rather than unlikely
Human fallback	Unrecognised replies reach a person instead of being dropped

BEYOND COMPLIANCE

The other way an agent fails: it goes *quiet*.

Compliance decides whether you are allowed to make the call. Architecture decides whether the call works when a vendor you do not control has a bad afternoon.

Your voice is a third-party dependency

A voice agent is not one product. It is a speech-to-text engine, a language model, a text-to-speech provider and a telephony carrier, stitched together in real time. **The text-to-speech layer is the one most people never think about** — and it is a separate company with its own status page. When it degrades, your agent does not return an error. It answers the phone and produces garbled audio, or nothing at all, to a real customer.

THE MITIGATION, AND IT TAKES ONE FIELD

Retell supports `fallback_voice_ids` — an ordered list of backup voices used automatically when the primary provider is having an outage. **The platform requires each fallback to be from a different vendor**, which is the entire point: a backup voice from the same provider fails at exactly the same moment as the primary one.

It is configured through the API rather than set once by hand, which is why it tends to be missing on builds that were clicked together in a dashboard. On our builds it is set on day one, not added after an incident.

The same principle, three more places

Silent failures	An integration that returns success while doing nothing is worse than one that crashes. Every write is read back — a status code is not evidence.
Visibility	Every workflow has an error path that alerts a human with the affected record. The default behaviour in most automation tools is to continue quietly.
Idempotency	A retry cannot distinguish *"the request failed"* from *"it succeeded and the response was lost"*. Every outbound action is keyed so that repeating it is safe.

These are not hypotheticals. Each one is a failure we have hit and documented on our own builds — including a reminder that delivered while its ledger row silently failed to write, which would have re-sent the same message on every subsequent cycle.

WHERE THIS LEAVES YOU

Three questions worth answering before the first dial.

1. **What does your opt-in actually say, word for word?** Not what it means — what it says. This decides whether an AI voice may call the numbers you already hold.
2. **Where do the leads come from?** Your own form is the clean case. Anything sourced elsewhere needs its consent evidence, and an indemnity in the contract.
3. **Which states are you calling into?** Disclosure and recording rules stack on top of the federal ones, and they are prompt-level fixes when handled up front.

Twenty minutes of decisions, baked into the schema on day one. They are the difference between a system that is fast and one that is fast *and* defensible.

TALK TO ONE

The fastest way to judge this work is to use it. Our demo receptionist is live and takes calls with no form and no signup — ask it to book you a viewing.

+1 (570) 673-0690

hisora.ai

On sources. Every claim in this document is drawn from the statute, the FCC rule or the docket cited beside it — 47 U.S.C. § 227, 47 CFR 64.1200, FCC Declaratory Ruling 24-17, and the decisions named on page 2. Nothing here is summarised from secondary commentary.

On what this is not. This is a working engineering reference, written to make build decisions with. It is not legal advice, and it does not replace counsel who knows your jurisdiction and your list.