2026 EDITION · FREE MARKET INSIGHT REPORT

The AI Communications Market Report

A practical map of AI voice and messaging for revenue teams: what changed, how the market is organized, the playbooks that convert, and how to deploy without joining the ninety five percent of pilots that never pay back.

16 SECTIONS | THE MARKET MAP | THE OPERATING STACK | PLAYBOOKS AND COMPLIANCE

01 · EXECUTIVE SUMMARY

Seven findings for busy readers

- **Speed is the cheapest conversion lever most teams have not pulled.** The two studies that measured it put qualifying odds around twenty one times higher at five minutes than at thirty, and roughly seven times higher inside the first hour than an hour later. Both were funded by companies selling software for the problem, which is worth knowing and does not make them wrong.
- **Voice quality stopped being the differentiator.** Every serious platform now delivers sub second, natural, interruptible conversation. Choosing by demo call is choosing by the part that no longer varies.
- **Every deployment runs a six layer stack.** The conversation, the runtime, telephony, orchestration, integrations, and operations. The demo shows one layer; the other five decide the outcome and are owned by someone, on purpose or by accident.
- **The market sorts into five categories, not fifty vendors.** Voice runtimes, no code builders, enterprise conversation platforms, point solutions, and managed layers. Naming your category first shrinks the shortlist to two or three real candidates.
- **Most of the statistics in this category do not survive checking.** The famous SMS and follow up numbers trace back to a retracted whitepaper and a 1942 survey of forty people. Section three shows the working, because a business case built on those numbers collapses the moment somebody checks.
- **The operational layer, not the model, is where pilots die.** A widely cited MIT study found ninety five percent of generative AI pilots show no measurable financial return within six months. The consistent cause: nobody owns tuning, monitoring, and upkeep after launch.
- **Compliance is a market entry ticket, not paperwork.** The United States prices consent mistakes per call, Europe prices them as a share of revenue, and India decides whether your calls connect at all. The rules are learnable, and section nine maps them.

21x

Higher qualification odds responding in 5 minutes vs 30, per widely cited lead response research.

95%

Of generative AI pilots show no measurable return within six months, per a widely cited MIT study.

6

Layers in every production deployment, whoever runs them. The demo shows you one.

---

02 · THE PROBLEM

The revenue leak, in numbers

Most teams do not lose deals in the negotiation. They lose them in the gaps: the form filled on Tuesday that gets a reply on Friday, the missed call that never gets a callback, the warm prospect whose follow up never went out. Each gap looks small. Together they are usually the largest untapped line in the funnel.

21x

Odds of qualifying a web lead called at five minutes rather than thirty. Source 1.

42hrs

Average time to respond to a web lead, among the companies that responded at all. Source 2.

~50%

Of companies audited with a test lead never responded to it at all. Source 2.

WHAT A DELAY COSTS, BY THE TWO STUDIES THAT MEASURED IT

Called within 5 minutes rather than 30 — **21x**

odds of qualifying the lead, InsideSales.com and Oldroyd, 2007

Called within the hour rather than an hour later — **7x**

odds of qualifying, separate study of 1.25 million leads, Harvard Business Review, 2011

Called after 24 hours — **60x worse**

against contact inside the first hour, same 2011 research

Both findings come from research funded by companies selling speed to lead software, and neither was peer reviewed. The 2007 study covered six unnamed companies; the 2011 work covered 1.25 million leads across 42 companies. They are the best evidence available, which is worth knowing before quoting them. Sources 1 and 2.

Two things drive this, and neither is fixed by hiring another person. The first is that buyer interest decays on the order of minutes. Someone filling in a form is comparing vendors at that moment, and every study above measures the same effect from a different angle: the gap between reaching them now and reaching them later is not incremental.

The second is that demand ignores office hours. A buyer researching at 9pm on a Sunday gets whoever answers, and voicemail loses to the next name on the list. Neither problem is a motivation problem. Both are scheduling problems, which is a much easier thing to solve.

> **The leak is a clock and a memory problem. Software is good at clocks and memory.**
>
> WHY THIS MARKET EXISTS

---

03 · THE EVIDENCE

The numbers this market repeats

Before we go further, a caution about the statistics in this category, including some we have quoted ourselves in the past. While preparing this report we traced the most repeated figures in AI sales and messaging back to their primary sources. Several do not survive the trip.

This is not pedantry. If you are building a business case on a number that turns out to be a 1942 survey of forty people, the case falls over in the meeting where someone checks.

THE CLAIM	VERDICT	WHAT THE SOURCE ACTUALLY SUPPORTS
SMS has a 98% open rate, and 90% are read within minutes	FOLKLORE	Traced to a 2010 industry whitepaper whose own authors have since walked it back, now estimating that roughly 55% of messages are actually read. SMS carries no read receipt infrastructure, so no open rate is measured the way email opens are.
80% of sales close between the fifth and twelfth touch, and 44% of reps quit after two	FOLKLORE	The only source ever cited is an unpublished 1942 survey of fewer than forty members of one local trade association chapter. The successor organisation publishes this history itself, as a warning. The 44% figure has no traceable source at all.
A text sent minutes before a call lifts pickup rates four times	FOLKLORE	The 4x figure comes from a vendor blog with no published method. The closest real evidence is a peer reviewed randomised trial in survey research, where pre notification texts raised response rates by roughly 60%. A real effect, an order of magnitude smaller than the claim.
Half of all sales go to the vendor who responds first	FOLKLORE	We searched the full text of both documents usually credited with this. It appears in neither. The figure circulates purely by mutual citation.
WhatsApp reaches 95% of smartphone users in India	OVERSTATED	Meta has disclosed 500 million plus users in India. Independent modelling puts it near 535 million, about 85% of the country's internet users, which is a narrower base than smartphone owners.
Business email open rates are 20 to 25%	REAL, WITH CAVEATS	True for genuine human opens, matching benchmarks from before 2021. Raw dashboard numbers now read 35 to 45% because Apple Mail Privacy Protection pre loads images. Quote one or the other, never both as if they measure the same thing.

THE CLAIM	VERDICT	WHAT THE SOURCE ACTUALLY SUPPORTS
95% of generative AI pilots show no measurable return	REAL, WITH CAVEATS	From a July 2025 MIT Media Lab NANDA report. It is preliminary and not peer reviewed, drawn from 52 interviews and a 153 person survey, and it counts only direct profit and loss impact within six months as success. Directionally useful, frequently overstated.
A new sales hire is productive in 90 days	REAL, WITH CAVEATS	Accurate for prospecting roles, where recurring industry surveys put ramp near three months. Full cycle reps carrying quota take far longer, with studies reporting nine months or more.

The pattern behind these is worth understanding, because it will keep producing new ones. A vendor publishes a figure that flatters its product. A blog quotes the vendor. A second blog quotes the blog and rounds the number up. Within a few years the original is unreachable and the figure has become something everybody knows. Speed to lead research is the clearest case: the two studies below are real and worth reading, but both were funded by companies selling software to fix the problem they measured, and the older one is now nineteen years old.

None of this means speed and follow up do not matter. Our own operating experience says they matter enormously. It means the honest version of the argument rests on mechanism rather than on a decorated statistic, and a vendor who cannot tell you where a number came from is quoting someone who could not either.

---

04 · THE TECHNOLOGY

What changed in AI voice, and why the demo is the wrong test

You could always tell. There was a beat before every reply, the delivery was flat, and anything off script produced the same polite non-answer. That has mostly stopped being true. Agents now answer inside a second, cope with being cut off halfway through a word, and stay with a buyer who switches language or doubles back on something said a minute ago.

<1s

Voice to voice response time production stacks now target. Slower reads as robotic within two exchanges.

50+

Languages and variants current speech models hold conversations in, including mid call switching.

2yrs

Roughly how long since voice quality stopped separating serious platforms.

That creates an awkward problem for anyone running an evaluation. **If every credible vendor sounds good, the demo has stopped being a test.** It shows you what the category can do. It tells you very little about what this particular vendor will do for you in a year. The questions that still separate them are the ones a demo never reaches:

- **Nobody picks up.** Does it recognise voicemail, leave something worth returning, text straight away, and try again at an hour when the person might actually be free?
- **The gap between touches.** Do SMS, WhatsApp and email share one memory with the call, or are they four tools that happen to be aimed at the same person?
- **The record.** Does the CRM update itself while the conversation is happening, or does it wait for notes that nobody writes?
- **Month six.** Is somebody reading transcripts and adjusting, or has the agent been running unattended since launch while your offer moved on without it?

Most of what follows is really about those four questions.

---

05 · THE FRAMEWORK

The operating stack: six layers, whoever runs them

Underneath the branding, every production deployment runs the same six layers. Which makes the useful evaluation question quite a narrow one: **which of these does the vendor run, and which quietly become your job?**

The conversation
A natural voice that answers instantly, in your buyer's language

WHAT BUYERS EXPERIENCE

Voice runtime
Speech recognition, the language model, synthesis, turn taking

RENTED, THEN CONFIGURED

Telephony and numbers
Carriers, caller ID, number reputation, regional deliverability
FAILS QUIETLY IF UNOWNED

Orchestration
Channel sequencing, shared memory, retries, provider failover
THE LARGEST HIDDEN BUILD

Integrations
CRM write back, calendars, webhooks, schema changes
BREAKS ON A TUESDAY

Operations
Monitoring, evals, prompt tuning, compliance updates
DECIDES MONTH SIX

Three of them cause most of the trouble. **Telephony** breaks without telling you: nothing throws an error, connect rates just sag as your numbers age and carriers start treating them differently. **Orchestration** is the biggest piece of invisible engineering on the list, and the one teams consistently underestimate, because queues and retry logic and shared memory are unglamorous work that only shows up when it is missing. Then there is **operations**, which is not a task at all. Prompts drift as your offer changes, models update underneath you, and the rules move in whichever markets you call.

Whoever owns those three owns the result. Every category in the next section is a different answer to that one question, dressed up in different marketing. There is a longer version of this framework in the production voice AI stack.

---

06 · THE MARKET MAP

Five categories, not fifty vendors

The landscape looks crowded until you sort it by operating model, at which point it collapses into five groups. Each one is a different split of those six layers between vendor and customer. Work out which group you are actually shopping in and the shortlist tends to pick itself.

CATEGORY 1 · VOICE RUNTIMES

Retell AI, Vapi, Bland AI, ElevenLabs Agents

WHAT THEY ARE

Developer platforms for the runtime layer: speech recognition, language model, synthesis, telephony, and function calling behind an API and SDK. The engines much of the market, including managed vendors, builds on.

GREAT WHEN

You have engineers who will own the system, want deep control, and treat voice as product. A strong team ships a working agent in days; the platforms are genuinely good.

WATCH FOR

The other five layers are yours: orchestration, integrations upkeep, number health, tuning, and compliance. The demo takes a weekend; the operating job is permanent.

INSIDE THE SET

Philosophy, not quality. Some package opinionated defaults for speed; some open every model choice for control; some center high volume dialing; one leads on raw voice quality. Our operator guides compare them directly: Retell AI vs Vapi and Retell AI vs Bland AI.

CATEGORY 2 · NO CODE BUILDERS

Synthflow and similar visual builders

WHAT THEY ARE

Visual flow builders that let an operator assemble a voice agent without engineers: templates, drag and drop logic, self serve onboarding.

GREAT WHEN

A capable operator wants to experiment fast, volumes are modest, and the motion is contained enough for a template to fit.

WATCH FOR

Self serve setup is not self running. Prompts, integrations, cross channel memory, and week by week tuning remain the customer job after the builder session ends.

CATEGORY 3 · ENTERPRISE CONVERSATION PLATFORMS

PolyAI, Yellow.ai, PKSHA (Japan)

WHAT THEY ARE

Enterprise grade conversational AI for contact centers and CX organizations: professional services led deployments, deep IVR replacement, procurement ready postures.

GREAT WHEN

You run enterprise inbound volume, need contact center scale, and have the timeline and organization for an enterprise engagement. The best of these are excellent at their job.

WATCH FOR

Altitude and motion fit. These platforms are strongest on inbound service resolution; outbound revenue motions and mid market speed are usually not the design center.

CATEGORY 4 · POINT SOLUTIONS — AI demo agents, schedulers, single job tools

WHAT THEY ARE	Focused products that do one conversational job well: a live software demo on a call, a booking flow, a single channel assistant.
GREAT WHEN	That one job is your bottleneck and the rest of your motion is already handled. A sharp tool beats a platform you use at five percent.
WATCH FOR	The moments before and after the tool job: the follow up, the CRM record, the second touch. Those live in the rest of your stack.

CATEGORY 5 · MANAGED LAYERS — Kaigen Labs and other managed operators

WHAT THEY ARE	Teams that run the whole stack on your behalf: they build on the runtimes from category one, add orchestration and integrations, and operate the system, tuning, monitoring, and compliance included.
GREAT WHEN	You want the outcome without staffing the operation, your motion spans channels, and speed to a measured pilot matters more than raw control.
WATCH FOR	Less direct control by design, and quality varies with the operator. Ask any managed vendor, including us, to show the operating loop: who reads transcripts, what gets tuned weekly, what you see in the dashboard.

How to use the map: pick the category that matches who will own layers three through six of the stack. Only then compare vendors inside it. Cross category comparisons are ownership decisions wearing a feature list.

---

07 · THE PLAYERS

Who is who, and what each one leaves you

The names a shortlist usually contains, sorted by the categories above. Every row describes a published product model, not a verdict: these are all real products with real deployments, and the last column is the honest question a demo will not answer.

VENDOR	WHAT IT IS	FITS	YOU STILL OPERATE
Vapi VOICE RUNTIME	Developer platform for voice agents: bring your own speech, language model, and voice, with the telephony handled.	Engineering teams that want to choose and swap every component.	Model selection, orchestration, integrations, tuning, compliance.
Bland AI VOICE RUNTIME	Enterprise voice platform running its own in house voice models, with a no code builder and a forward deployed engineering path.	Regulated and high volume operations that want to own the calling machine.	Orchestration, integrations, tuning, compliance.
Retell AI VOICE RUNTIME	Voice agent platform with a proprietary turn taking layer, a low code builder, and a separate chat agent product.	Teams that want the fastest path to a reliable production agent.	Cross channel sequencing, integrations upkeep, tuning, compliance.
ElevenLabs Agents VOICE RUNTIME	Agent runtime on top of the company's own voice models, with a visual builder and an enterprise tier.	Deployments where voice quality or a cloned brand voice decides the evaluation.	The sequence itself: what happens between touches, and the tuning after launch.
Synthflow NO CODE BUILDER	Visual flow builder for voice agents, with APIs underneath and a growing enterprise motion.	Operators who want to launch without a development team.	Prompts, flows, integrations, and the weekly tuning after launch.
Supersonik POINT SOLUTION	AI agent that joins a live video call and runs a guided demo of your own software.	Software teams whose bottleneck is demo capacity.	Everything around the demo: follow up, CRM, the wider motion.

VENDOR	WHAT IT IS	FITS	YOU STILL OPERATE
PKSHA Technology ENTERPRISE PLATFORM	Japanese enterprise AI whose ChatAgent and VoiceAgent products serve domestic contact centre and communication workloads.	Japanese enterprises buying domestic AI with local support.	Sales motion design outside the contact centre remit.
PolyAI ENTERPRISE PLATFORM	Enterprise voice agents built on its own speech model, deployed with forward deployed engineers and dialogue designers.	Enterprise inbound volume and deep IVR replacement.	Outbound sequencing and the sales motion beyond service.
Yellow.ai ENTERPRISE PLATFORM	Enterprise CX automation across chat, voice, and email, sold with systems integrator partners.	Large support organisations automating service conversations.	Sales specific orchestration and pipeline write back.
Kaigen Labs MANAGED LAYER	Managed layer that runs voice, WhatsApp, SMS, and email as one sequenced motion on top of the runtimes above.	Mid market revenue teams that want the outcome without staffing the operation.	The judgment calls. We run the rest.
In house build IN HOUSE	Your engineers assemble the stack on a runtime and operate all six layers.	Companies where conversation is core product, with the team to staff it.	Every layer, permanently.

Two layers sit underneath all of them and are worth naming, because they are often mistaken for competitors: **telephony and messaging carriers** such as Twilio, and **realtime media infrastructure** such as LiveKit. Almost every product above rides on

one or both. If a vendor cannot tell you which carrier and which media layer they run on, that is a useful question to keep asking.

---

08 · THE POSITIONING MAP

Where each player sits

Two axes, both structural. Across: **who operates the system once it is live**, from a platform you run yourself to a service run for you. Up: **how much of the motion the product covers**, from the phone call alone to voice and messaging on one shared memory.

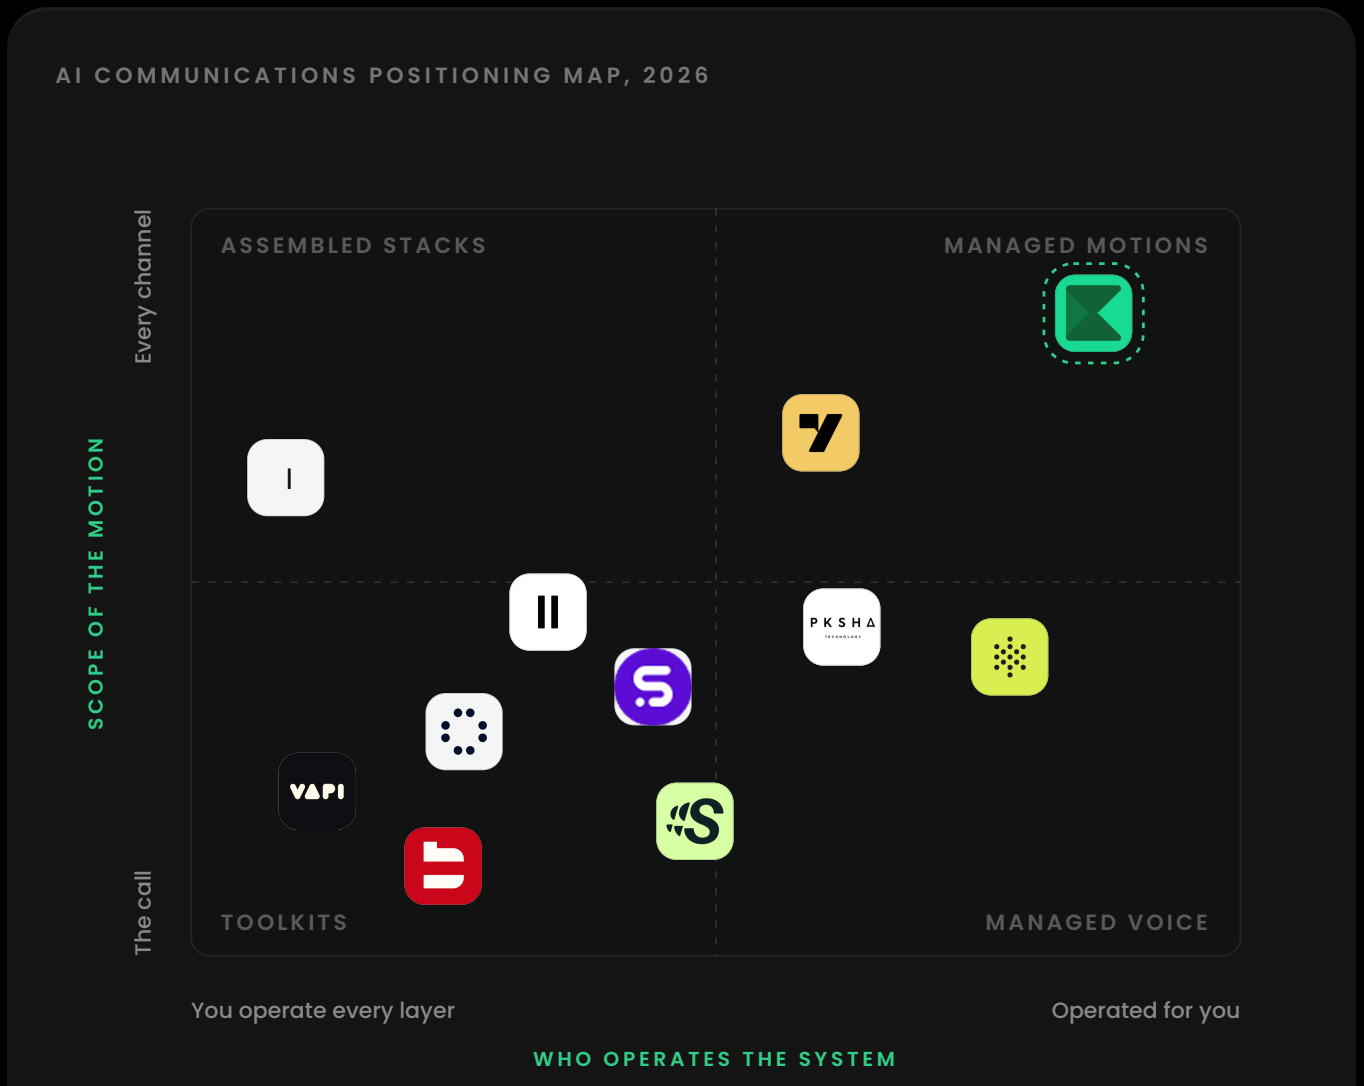

Vapi
Voice runtime

Bland AI
Voice runtime

Retell AI
Voice runtime

ElevenLabs Agents
Voice runtime

Synthflow
No code builder

Supersonik
Point solution

PKSHA Technology
Enterprise platform

PolyAI
Enterprise platform

Yellow.ai
Enterprise platform

Kaigen Labs
Managed layer · that is us

In house build
In house

Placement reflects each vendor's published product model, not quality, market share, revenue, or customer satisfaction. No vendor was scored, surveyed, ranked, or paid to appear. Kaigen Labs publishes this report and is plotted on it, marked accordingly.

How to read it

The top right is not a winner's circle. It is the highest service, lowest control corner of the market, and plenty of teams belong in the bottom left instead. A company building a voice product should be suspicious of anything that takes the runtime out of its hands; a ten person sales team with no engineers should be equally suspicious of a toolkit. The map tells you which trade you are making, not which vendor is better.

- **Toolkits, bottom left.** Voice runtimes and in house builds. The product is the call, and every layer around it is yours. Cheapest per minute, highest total ownership.
- **Assembled stacks, top left.** Multi channel motions that exist because a team built the orchestration. Powerful and specific to you, and the maintenance never ends.
- **Managed voice, bottom right.** Enterprise platforms and point solutions where someone else runs a scoped job well. Strong inside the scope, quiet outside it.
- **Managed motions, top right.** Operators who run the whole stack across channels. You buy an outcome and give up direct control of the components.

Three honest caveats. Vendors move: several runtimes are adding managed tiers, and several enterprise platforms are adding outbound. Placement is by primary product model, so a vendor with a services arm sits where the bulk of its customers actually land. And a dot near an axis is a judgment call about emphasis, not a measurement.

---

09 · THE PLAYBOOK

What converts: speed, then sequence

Across industries, the deployments that pay back share two mechanics. First, **instant response**: the agent calls or replies within a minute of the trigger, while intent is at its peak. Second, **a sequenced motion**: the call is one touch in a coordinated cadence

across text, voice, and email, all sharing one memory of the buyer.

Section three is the reason this one is thinner on statistics than you might expect. The famous channel numbers, the ninety eight percent SMS open rate and the four times pickup lift from a pre call text, did not survive checking, so we are not going to repeat them here to make the case look stronger. What follows rests on mechanism and on what we see running these systems.

The mechanism is straightforward. A buyer who has been told a call is coming answers a number they would otherwise ignore, because the call arrives with context instead of as an interruption. The nearest rigorous evidence for that effect comes from survey research rather than sales: a randomised trial found pre notification texts raised response rates by roughly sixty percent. Real, useful, and a long way short of the four times a vendor blog claimed.

A reference cadence

Adapt days and channels to your market; in India, WhatsApp typically takes the lead role.

Day 1
Text pre warm
"Our assistant will call tomorrow about your enquiry."

Day 2
Text + AI call
Five minute heads up, then the call. Voicemail and instant text if no pickup.

Day 4
Email
The document discussed, or a relevant one pager.

Day 5
Call retry
Different hour. References the email, not a cold rerun.

Day 7
Polite breakup
"Reply whenever you are ready." Door stays open.

- **One memory across every touch.** The call references the text; the email references the voicemail. Disconnected tools reintroduce the leak the system exists to close.
- **Every touch lands in the CRM as it happens**, with outcome, objections, and next step, so a human can take over any conversation with the full thread in front of them.

- **Judgment escalates to people.** The system carries volume; pricing exceptions, relationships, and negotiations route to a human with context attached.

---

10 · INDUSTRY LENSES

The same leak, four different disguises

The mechanics travel, but they show up wearing local clothes, and the version you recognise is usually the one from your own industry. Four from markets where the conversation is most of the job:

RECRUITMENT AND STAFFING

The first agency to reach a candidate usually wins the placement.

Industry figures put the large majority of placements with the first responder, while recruiters spend most of their day on repetitive first pass screening. So the agent takes the first pass on every new application and the recruiters come in at the shortlist, which is the part they are good at.

WATCH time from application to first conversation

REAL ESTATE

Enquiries arrive at 9pm; viewings are lost by 9am.

Portal leads go cold overnight and no show rates quietly tax every diary. Answering within a minute at any hour is most of the fix; qualifying budget and timeline, booking the viewing and chasing the reschedule is the rest of it.

WATCH enquiry to booked viewing rate, no show rate

HOSPITALITY

The front desk cannot answer two phones at midnight.

After hours callers who reach voicemail book the next hotel on the list. Every booking call gets answered on the first ring, whatever the hour, and the upsell and pre arrival messages go out on whichever channel the guest already uses.

WATCH answered call rate after hours, direct booking share

SOFTWARE AND SAAS

High intent visitors ask layered questions; a chat widget collects emails.

Pricing page visitors and trial signups are comparing vendors in the moment. What works here is answering the whole stack of questions live and booking the demo inside the same conversation, with the CRM already updated by the time a rep opens their laptop.

WATCH speed to first touch, demo booked per high intent visit

If your industry is not listed, the transfer test is simple: find where interest decays fastest in your funnel, and put the always on layer exactly there first.

---

11 · THE DECISION

Build or buy: compare capacity, then ownership

Most build or buy cases get argued on cost per seat against cost per licence, which compares two things that do not produce the same output. Capacity is the more useful axis: what each option lets you do in a day, and how quickly you can turn it up. Then the second question, which is who ends up operating the stack.

DIMENSION	HUMAN TEAM	AI SYSTEM
Speed to first touch	Whenever a rep next works the list	Under a minute from the trigger
Coverage	Business hours, one or two time zones	Around the clock, every market you sell into
Ramp	Months per hire, per sales onboarding benchmarks	Weeks to a live pilot
Consistency	Varies by rep, day, and morale	The same playbook on every conversation
Judgment and relationships	The entire reason to have people	Escalates to people, with context attached

TIME TO FIRST PRODUCTIVE OUTPUT

New prospecting hire reaches full productivity — **~3 months**

recurring industry survey of B2B software teams; full cycle reps take considerably longer

Managed pilot live and taking conversations — **2 to 3 weeks**

the Kaigen Method, and comparable managed engagements

Ramp figures come from recurring industry surveys of business software sales teams and describe prospecting roles; full cycle reps carrying quota take considerably longer. The pilot timeline is the Kaigen

Method and comparable managed engagements. Two different kinds of work, compared on the one axis a buyer feels.

TURNOVER IN PROSPECTING ROLES

Annual turnover in prospecting roles — **30 to 50%**

recurring industry compensation and metrics surveys

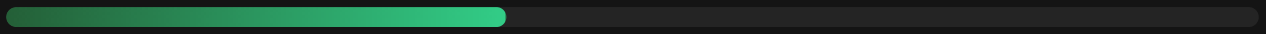

Annual turnover among account executives — **~30%**

for comparison, from the same body of research

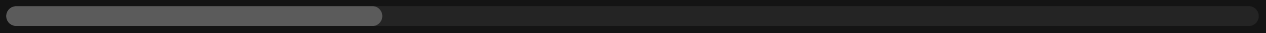

The best supported figure in this report, from recurring compensation and metrics surveys of business software sales organisations. Turnover is why capacity built from headcount alone is harder to hold than it looks on a plan.

The teams getting the strongest returns are not replacing people. They split the work: software carries first pass volume and follow up; people keep the conversations where judgment moves the number. Then the build or buy question reduces to four questions about the stack.

Q1

Who tunes prompts in month six?

A name on your payroll, a vendor operating team, or nobody. Nobody is how the ninety five percent happens.

Q2

Who owns number health and deliverability?

Reputation decays quietly with volume. Someone watches it weekly or pickup rates sink without an error message.

Q3

What happens when a provider has an outage?

Every provider has one eventually. Single runtime setups cannot fail over to themselves; ask for the plan B.

Q4

Who fixes the CRM sync when a field changes?

Schema changes are routine. If one breaks the pipeline for a week, the integration was a demo.

The rules that decide whether your calls connect

Every market runs on the same skeleton: get permission, respect the do not call lists, call at humane hours, say who you are, and stop when asked. The details differ, and the details are what regulators fine. Treat this as a working map rather than legal advice, and have counsel sign off before your first campaign.

MARKET	KEY RULE	WHAT IT MEANS FOR AI OUTREACH
United States	TCPA, plus the February 2024 FCC ruling that AI voices count as artificial voices	Prior express consent for marketing calls, do not call scrubbing, and calling windows under both FCC and FTC rules. Damages start at $500 per call and treble for willful violations, and there is a private right of action, so class actions compound quickly. A dedicated AI disclosure rule was proposed in 2024 but has not been adopted. Sources 11, 12, 13.
EU and UK	GDPR, with PECR alongside UK GDPR	A lawful basis before outreach and data rights on a clock. Fines reach 4% of global turnover or 20 million euro at the higher tier, and 2% or 10 million at the lower. In the UK, screen against the Telephone Preference Service, and note that automated calls have always required consent outright. Sources 14, 15.
India	TRAI commercial communication rules and the do not disturb registry	Register as a sender, use the designated 140 series for promotional traffic and the 1600 series for transactional, and scrub against the registry before dialing. Enforcement has tightened since 2024, with suspensions on a first violation. Get this wrong and the calls simply do not connect. Source 16.
Canada	Two regimes, and the distinction matters	CASL governs electronic messages such as email and SMS, with penalties up to $10 million per violation for a business. Voice calls are explicitly carved out of CASL and fall under the National Do Not Call List instead, where penalties are far lower. Reports that apply the CASL number to AI calling are wrong. Source 17.

MARKET	KEY RULE	WHAT IT MEANS FOR AI OUTREACH
Australia	Do Not Call Register Act and the telemarketing standard	Scrub the register, and keep to weekdays 9am to 8pm and Saturdays 9am to 5pm, with no calls on Sundays or national public holidays. The rules are technology neutral, so AI calls appear to be in scope, though the regulator has issued no AI specific ruling. Source 18.
WhatsApp	Meta platform policy	Opt in required, a 24 hour customer service window that resets on each user message, and pre approved templates outside it. The strictest gatekeeper of any channel here, and the one with the highest engagement. Source 19.

- **Disclosure is now table stakes for AI calls** in major markets: the agent says what it is at the start, every time.
- **Consent has a shape:** written, naming the business, tied to the number, never a condition of purchase. A clear checkbox does it; a bought list never does.
- **Rules move.** Whoever owns your operations layer owns tracking them, per market, as part of the standing job.

The market by market detail, including primary sources, is in our voice AI compliance guide.

---

13 · GETTING STARTED

The ninety day blueprint

Build it, buy a platform, or hire an operator: the deployments that pay back tend to move in the same order anyway. Find the leak. Build for that one leak. Put it in front of real volume while somebody is still watching. Then keep tuning it, which is the part most teams stop doing around week five.

PHASE	WHAT HAPPENS	LEAVE WITH
Weeks 1 to 2 · Map	Trace the funnel end to end: where leads enter, how fast they hear back, where conversations stall, where deals quietly die. Rank the leaks by revenue impact.	One page of your real leaks, ranked, with baseline numbers.

PHASE	WHAT HAPPENS	LEAVE WITH
Weeks 2 to 4 · Build small	One workflow, not a transformation: the single highest impact leak. Conversation playbook, channels, CRM wiring, compliance checks for your markets.	A working system on a focused slice, with a human in the loop.
Weeks 4 to 8 · Launch narrow	Real volume on the slice. Watch transcripts daily at first. Measure against the baseline from the map phase, not against a vendor deck.	Honest numbers: contact rate, booking rate, speed to first touch.
Weeks 8 to 13 · Tune and expand	Weekly loop: read transcripts, run evals, fix the worst objection handling, then widen the slice or add the next workflow.	A system improving week over week, and a decision you can defend.

The mapping phase is the one people want to skip, and it is the one that decides whether the rest means anything. Without a baseline from your own funnel you end up comparing results to a vendor's slide, which is not a comparison. Start wherever interest decays fastest, usually speed to lead, because that is where the difference between doing nothing and doing something is most visible. And whoever ends up running the system, ask how often it gets tuned. That cadence predicts month six better than anything in the original build.

> **Ninety days is enough to find out. The projects that fail rarely failed at the technology, they just stopped paying attention somewhere around week five.**
>
> THE KAIGEN TEAM

---

14 · ABOUT THE AUTHORS

Where Kaigen Labs sits in the map we drew

We should say plainly where we sit. **Kaigen Labs is a category five vendor**, a managed layer. We build on the runtimes from category one, pick between them per deployment, and then run the rest of the stack for mid market revenue teams: voice, SMS, WhatsApp and email as one motion, CRM write back included, monitored and tuned every week, with most pilots taking calls inside two to three weeks.

None of this report is written from the outside. The stack in section five is our own architecture map, the cadence in section nine is a motion we actually run, and the four questions in section eleven are the ones we tell prospects to put to us.

- **If you are building:** our operator guides compare the runtimes honestly, including Retell AI vs Vapi, Retell AI vs Bland AI, and the full production stack.
- **If you are evaluating the managed path:** we publish side by side comparisons against every major alternative, half discs and all.
- **If you want the map drawn on your funnel:** a twenty minute call with the Kaigen team. You bring how you sell; we sketch the layers for your case and say honestly which ones you could run yourself.

HEAR IT, THEN JUDGE IT

A live Kaigen voice agent runs in your browser

Pick an industry, start a call, and ask whatever you want. Or bring your funnel to a twenty minute audit and we will map the six layers against how you sell today.

---

15 · METHOD AND SOURCES

How this report was put together

If a report is going to place other companies on a chart, it owes you its method. Here is ours, including the things we chose not to do.

- **None of the statistics are our own research.** Every figure is a widely cited industry finding, and each one carries its attribution in the sentence or the chart source line. Where a number is directionally useful but loosely sourced, we say so rather than dressing it up.
- **The positioning map plots product model, not performance.** The axes are who operates the system and how much of the motion the product covers, both readable from what each vendor publishes about itself. Nobody was scored, surveyed, ranked, or charged to appear, and no quadrant is labelled as the good one.
- **We are on the map.** Kaigen Labs publishes this report and sits in the managed motions quadrant, marked as ours in the legend. Read our placement with the scepticism you would apply to any vendor mapping itself.

- **We publish no market size figure on purpose.** Published estimates for this category differ by an order of magnitude depending on whether they count contact centre software, speech APIs, or agent platforms. A number that swings that far with its definition tells a buyer nothing useful.
- **Vendor descriptions come from two places:** what each company publishes about its own product, and our experience running several of these platforms in production. Where we have not operated something, we describe the category rather than inventing detail.
- **Corrections are welcome.** If you work at a company on this map and we have described your product wrongly, tell us and we will fix it in the next edition.

---

16 · REFERENCES

Where to check everything above

Every external claim in this report, with somewhere to verify it. Where a source has a commercial interest in its own finding, or a limitation worth knowing, that is noted rather than buried.

1 **Speed to lead: 21x qualifying odds at five minutes versus thirty, and 100x contact odds**

InsideSales.com with Dr James Oldroyd, Lead Response Management Study, 2007. Vendor funded, never peer reviewed.

content.marketingsherpa.com

2 **Response time distribution and the 7x and 60x qualifying findings**

Oldroyd, McElheran and Elkington, The Short Life of Online Sales Leads, Harvard Business Review, March 2011

hbr.org

3 **95% of generative AI pilots showing no measurable profit and loss impact**

MIT Media Lab Project NANDA, The GenAI Divide: State of AI in Business 2025, July 2025. Preliminary, not peer reviewed.

mlq.ai

4 **Enterprise AI projects abandoned after proof of concept**

Gartner press release, 30% of generative AI projects will be abandoned after proof of concept, July 2024

gartner.com

5 **Companies struggling to scale value from AI**

Boston Consulting Group, Where's the Value in AI?, October 2024

bcg.com

6 **The SMS open rate claim, and its retraction by the original researchers**

Mobilesquared, SMS engagement dropping but still outperforms other channels, August 2022

mobilesquared.co.uk

7 **The origin of the five to twelve touches claim: a 1942 survey of fewer than forty people**

Sales and Marketing Executives International, on the provenance of widely quoted sales statistics

smei.org

8 **Pre notification messages and response rates, the nearest peer reviewed evidence**

Bull et al., Pre-Survey Text Messages Improve Participation Rate in an Australian Mobile Telephone Survey, PLOS One, 2016

journals.plos.org

9 **WhatsApp reach in India**

DataReportal with Kepios, Digital 2025: India

datareportal.com

10 **Email open rate benchmarks, and the effect of Apple Mail Privacy Protection**

Mailchimp email marketing benchmarks, and Litmus research on Apple Mail Privacy Protection

mailchimp.com

11 **AI generated voices count as artificial voices under the TCPA**

Federal Communications Commission, Declaratory Ruling FCC 24-17, CG Docket 23-362, February 2024

docs.fcc.gov

12 **TCPA statutory damages of $500 per violation, trebled for willful violations**

47 U.S. Code section 227

law.cornell.edu

13 **The one to one consent rule vacated, a separate action from the AI voice ruling**

Insurance Marketing Coalition v. FCC, Eleventh Circuit, January 2025

media.ca11.uscourts.gov

14 **GDPR fine tiers: up to 4% or 20 million euro at the higher tier, 2% or 10 million at the lower**
Regulation (EU) 2016/679, Article 83
gdpr-info.eu

15 **UK marketing call rules and the Telephone Preference Service**
Information Commissioner's Office guidance on direct marketing and PECR
ico.org.uk

16 **India commercial communication rules, the 140 and 1600 series, and the do not disturb registry**
Telecom Regulatory Authority of India, TCCCPR 2018 and subsequent directions
trai.gov.in

17 **Canadian rules: CASL covers electronic messages, while voice calls fall under the National Do Not Call List**
CASL section 6(8), and CRTC Unsolicited Telecommunications Rules
laws-lois.justice.gc.ca

18 **Australian telemarketing calling hours**
Telecommunications (Telemarketing and Research Calls) Industry Standard 2017, ACMA
donotcall.gov.au

19 **WhatsApp 24 hour customer service window, template approval, and opt in**
Meta, WhatsApp Business Platform documentation and messaging policy
developers.facebook.com

20 **Market size estimates spanning roughly five times in the same year, by definition**
Compiled from published estimates by MarketsandMarkets, Precedence Research, and IMARC Group
marketsandmarkets.com

The AI Communications Market Report, 2026 edition. Written by the Kaigen team, verified in August 2026. Guidance, not legal or financial advice. You are welcome to share it in full.