ALLIED BizTech SOLUTIONS

FIELD GUIDE · 2026

Rebuild Gong yourself.
The plan we follow.

For sales teams of 30 to 300 reps who want to own their call recordings, transcripts, AI insights and coaching, and stop paying per seat for models they could call directly.

- → Is it worth it, at your size
- → What to rebuild and what to drop
- → The architecture, in three layers
- → Migration and go-live checklists
- → The data model
- → A week-by-week build plan
- → Transcripts and insights that hold up
- → Where teams get stuck

Gong charges per seat. *The models charge per minute.*

Gong bills per seat, about $200 a month on Pro and more bundled, plus a platform fee of $5K–$50K a year. Underneath sit Whisper at $0.006 a minute of audio and Claude at about a cent per 30-minute call. AI costs have dropped ~80% in 18 months; Gong's price has not. A rebuild pays the models directly.

TEAM	GONG A YEAR	REBUILT: YEAR ONE / AFTER	SAVED IN YEAR ONE
15 reps · standard plan	~$36,000	~$24,700 / $4,800	~$11K
30 reps · Gong Pro	~$72,000	~$25,900 / $6,000	~$46,100
50 reps · mid-stage B2B SaaS	~$120,000	~$30,700 / $10,800	~$89,300
200 reps · Premium + Forecast	~$540,000	~$87,800 / $48,000	~$452,200
500+ reps · custom contract	~$1,300,000+	~$169,800 / $130,000	~$1,130,200

Gong list prices, May 2026. "Rebuilt" runs on the Whisper API, Claude Sonnet for analysis ($3 input / $15 output per million tokens) and AWS S3 and EC2, plus our fixed build price: $19,900 up to 100 reps, $39,800 (the top of the $29,900–$39,800 scale tier) above 100. If your own team builds it, put your engineering cost in that place.

THE RULE OF THUMB

Even at 15 reps, year one saves ~$11K and $31K a year comes back after. **At 50 reps, the rebuild saves ~$89K in year one and ~$109K a year after.** From 200 reps, year one alone saves more than $450K.

Who you need

- **A senior full-stack engineer** at home with Postgres, S3, the Zoom, Meet and Teams recording APIs and an LLM API. Our team ships this in 5–7 weeks; one engineer working between other priorities should plan for longer.
- **Your sales enablement lead**, two hours a week: the sales method, the insights managers read, the five pilot reps, the go-live call.
- **An owner after launch**: hosting, monitoring, the monthly model and prompt review, consent rules, and on-call.

Rebuild the 15–25% you use. *Drop the rest.*

Gong sells a platform. A sales team uses recording, transcripts, a handful of insights, coaching and the Salesforce sync every day. Those are the project; the proprietary extras are where rebuilds stall.

REBUILD ONE FOR ONE

- **Auto-recording** from Zoom, Meet and Teams: passive, no change to the rep's day
- **Phone call recording** through Twilio, stored in your S3
- **Transcripts** with speaker diarization, timestamped and searchable
- **Talk ratio and sentiment** per call, per rep, per deal
- **Objection, topic and question detection**, tuned to your sales method
- **Deal warnings**: missing topics, quiet champion, stalled deals
- **Coaching**: 1:1 prep, highlight clips, scorecards, a manager review queue
- **Call library**: search, filter, clip-and-share
- **Salesforce sync** (activity log, call notes, deal insights) and **Slack or email alerts**

REBUILD, SIMPLER

- **Deal-risk scoring**: rules (no champion identified, single thread, no manager mention) in place of Gong's ML
- **Forecast**: a deal and pipeline forecast on your Salesforce data, as an add-on; or Clari or Aviso direct

DROP ON PURPOSE

- **Gong Engage**: sequencing overlaps with Outreach; rebuild that on its own
- **Industry benchmarks**: they exist only on Gong's pooled customer data
- **Battlecard auto-generation**: Crayon or Klue
- **The mobile coaching app**: stay on the web; a mobile rebuild rarely pays back

ADD WHAT PER-SEAT PRICING DISCOURAGES

- New reps at no seat fee; cost follows call volume
- Insight prompts written for your sales method
- SQL access to every transcript and insight, from your warehouse

Three layers. *All standard parts.*

Conversation intelligence is three standard capabilities: speech-to-text, a language model for analysis, and storage for recordings and transcripts. Gong wraps them; you call them directly.

LAYER 1 · CAPTURE

Recording and transcripts

- Zoom, Meet and Teams webhooks
- Phone calls via Twilio recording
- Whisper API direct
- Speaker diarization, timestamps
- S3, retention you set

LAYER 2 · ANALYSIS

LLM insights

- Claude or OpenAI direct
- Talk ratio, sentiment, objections
- Deal warnings, coaching topics
- Prompts for your sales method
- Results in Postgres for reporting

LAYER 3 · UI

The workspace

- Call library, clip-and-share
- Coaching dashboard, 1:1 prep
- Deal intelligence, pipeline risk
- Salesforce sync
- Slack and email alerts

Decisions to make in week one

DECISION	OUR DEFAULT	WHY
Where it runs	Your AWS account	Recordings and transcripts never leave your tenancy
Capture	Platform recording webhooks; a bot only where needed	Nothing changes for the rep
Transcription	Whisper direct; Deepgram for real time	The same class of model, billed by the minute
Analysis	One model provider behind a small adapter	Swap models without touching the workspace
CRM	The same Salesforce fields Gong writes today	Salesforce workflows keep working on day one
Sign-in	Your SSO from day one	Offboarding happens where it already does

Compliance: handle consent at capture (two-party consent, EU recording laws, sector rules) with the same disclaimers Gong uses. On your own AWS account, a SOC 2 and HIPAA-ready posture is yours to hold.

Ten tables *carry the whole product.*

Model these before writing any screen. Keep Gong's ID on every imported call, so the import can run twice and any insight traces back to its source.

TABLE	KEY COLUMNS
`calls`	id, external_id, source (zoom, meet, teams, phone), started_at, duration, rep_id, opportunity_id, account_id, recording_key, consent_status, language
`participants`	call_id, email, name, side (internal, external), speaker_label, talk_seconds, crm_contact_id
`utterances`	call_id, speaker_label, start_ms, end_ms, text, confidence; a full-text index across all calls
`insights`	id, call_id, type (objection, topic, question, next step, competitor), text, start_ms, model, prompt_version
`call_metrics`	call_id, talk_ratio, longest_monologue, questions_asked, sentiment
`deal_warnings`	opportunity_id, rule (no champion, single thread, no manager mention, stalled), raised_at, cleared_at
`scorecards`	id, call_id, reviewer_id, template_id, scores (jsonb), comments, shared_at
`clips`	id, call_id, start_ms, end_ms, created_by, note, shared_with
`prompts`	id, name, methodology (MEDDIC, Challenger, your own), version, text, active
`model_runs`	call_id, step (transcribe, analyze), model, audio minutes or tokens, cost, status, created_at

TWO RULES THAT SAVE A REBUILD

Keep the raw transcript and the prompt version on every insight: when a prompt improves, re-run old calls and compare. **Log every model call with its minutes, tokens and cost**: it is how you prove the saving and catch a runaway prompt.

Seven weeks, *pilot included.*

This is the order we build in. Each week ends with something your sales enablement lead can try, and the last two are a pilot with five real reps.

WEEK	BUILD	DONE WHEN
1	**Export and model.** The ten tables. Bulk-export the last 12 months of recordings, transcripts and insights from Gong's API into S3 and staging. Reps mapped to Salesforce users. SSO.	Counts match Gong; 20 calls spot-checked end to end
2	**Capture.** Recording webhooks for Zoom, Meet and Teams, Twilio phone recording, consent at capture, S3 with retention.	A rep's calls record with no change to their day
3	**Transcripts.** Whisper with diarization and timestamps, speakers matched to people, full-text search across every call.	Test calls transcribe with every speaker named right
4	**Insights.** Prompts for your sales method, talk ratio, sentiment, objections, deal-warning rules, tested against Gong (next page).	Managers rate the insights as good as Gong's
5	**Workspace and Salesforce.** Call library with clips, coaching dashboard, 1:1 prep, scorecards, Salesforce logging, Slack alerts.	Every call shows on the right opportunity
6–7	**Pilot and cutover.** Five reps for two weeks, a daily fix list, a delta export from Gong, then the recorder on for every rep and Gong's off.	Managers coach from it, not from Gong

AFTER CUTOVER

Switch Gong's recorder off the same day, so no call records twice. Keep the account **read-only until your renewal date**, then cancel. Older calls stay one click away while the team settles in.

Clean transcripts in. *Sharp insights out.*

Gong runs on foundation speech models; Whisper direct gives the same accuracy, diarization and languages. The quality of what managers read comes from three things you control: clean audio, the right speakers, and prompts written for your sales method.

1. **Capture clean audio.** Use the platform's own recording where you can, with a track per speaker where it offers one; record phone calls in two channels. Clean audio fixes more errors than any model change.
2. **Name the speakers.** Diarization finds the voices; your calendar and CRM say who they are. Match each speaker to a rep or a contact, so talk ratio and objections land on the right person.
3. **Teach it your words.** Pass your product, feature and competitor names to Whisper as a prompt. Those are the words insights depend on.
4. **Compute what you can, prompt the rest.** Talk ratio, monologues and questions come from timestamps, not a model. Use Claude for objections, topics, next steps and sentiment, with structured output and a timestamp for each finding, so every insight opens at the moment in the call.
5. **Test before launch.** Take recent calls your managers know well, run them through and put each insight next to Gong's. Launch when your enablement lead signs off; each miss points at a prompt to fix.

Your method, not a generic one

Write the prompts for MEDDIC, Challenger or your own playbook, about your product and buyers. Insights tuned to your sales motion are often sharper than Gong's generic ones.

Measure the saving

Track audio minutes, tokens and cost per call next to the Gong bill you replaced, with the share of insights managers act on. Review the misses monthly.

Move it, launch it, *run it.*

Migration

- ☐ Export the last 12 months: recordings, transcripts, insights, highlights
- ☐ Export twice: everything in week one, the changes since at cutover
- ☐ Keep Gong's call ID on every row
- ☐ Map Gong users to reps and Salesforce users before importing
- ☐ Link every call to its opportunity, contact and account
- ☐ Recreate scorecard templates and keyword trackers
- ☐ Copy recordings to S3 with your retention rules
- ☐ Check counts, then play and read 20 calls end to end

Go-live

- ☐ The recorder on for every rep's calendar and phone line
- ☐ Consent notices live for every platform and region
- ☐ Reps and managers signing in through your SSO
- ☐ Salesforce logging on the fields Gong writes
- ☐ Alerts on failed recordings, transcription backlog and sync errors
- ☐ Gong's recorder off the same day, so no call records twice

Running it

- ☐ **Always on:** capture and transcription error alerts, security patches
- ☐ **Weekly:** deal warnings and coaching activity, to sales leaders
- ☐ **Monthly:** model upgrades, a prompt review against flagged misses, vocabulary updates
- ☐ **Quarterly:** cost review, model and AWS usage against the Gong bill you replaced; retention review
- ☐ **Yearly:** a restore drill, an access review and a consent-rules review

Who owns it

Name one person before launch. Every sales call passes through it, with your customers' voices in it: someone has to answer when the recorder misses a big demo or a consent rule changes.

Five ways a rebuild *stalls.*

1. Switching vendors and calling it a saving

Chorus, now ZoomInfo Conversation Intelligence, prices per seat much like Gong: the same economics under another brand. Salesloft Conversations bundles sequencing with intelligence at similar per-seat economics; it suits teams already on Salesloft.

2. A cheaper per-seat tool

Wingman, Avoma and Fathom cost far less at small seat counts, $30–80 a seat a month. The per-seat question comes back at scale.

3. "Just call the APIs"

You can. Transcripts are the easy part, and a $15K freelance build gets you those and a basic dashboard. Insight prompts, deal warnings, coaching, Salesforce sync and the call library are real engineering.

4. Chasing Gong's machine learning

Gong's deal-risk model is trained on millions of calls across its customers; that is its moat. Rule-based flagging covers 80% of what most teams use it for.

5. Skipping the pilot

Five reps for two weeks find what a test plan misses: the competitor name the transcripts get wrong, the scorecard a manager fills in every Friday.

And one that is not a mistake: negotiating

Gong discounts Pro and Enterprise 15–25% on multi-year, multi-hundred-seat commits. Do it if the rebuild is a year away; it captures less than a fifth of what a rebuild does.

RATHER HAVE IT BUILT?

We build this exact plan. *Fixed price.*

The same scope, architecture and seven weeks, delivered by a senior team that has shipped 250+ products.

$19,900 fixed up to 100 reps: a senior engineer, a coder and a half-time delivery manager for 6 weeks; $29,900–$39,800 from 100 to 300; agreed before we start

Live in 5–7 weeks, with a two-week pilot with five of your reps first

Migration included: the last 12 months of your Gong call library, coaching highlights included

Forecast as an add-on: deal and pipeline forecasting on your insights, +$6,600–$11,600

90-day warranty; the code in your GitHub, recordings and transcripts in your AWS

We can run it for $3,600 a year: monitoring and monthly model upgrades

alliedbiz.tech/rebuild/gong-alternative

Allied BizTech · 15+ years · ISO 9001:2015 · 250+ products shipped. Gong is a trademark of its owner, named here for identification only; this guide is independent and not endorsed by Gong. Prices as of May 2026.