

Selling system

SEND NO

NEVER FULLER

DEMO DID +1 (650) 476-2005 ONLY

WEB RENDER · LEADS REDACTED

AT A GLANCE

How a stranger becomes a paying client: 6 stages, the 4-offering menu, gap analysis G1-G10.

Updated 2026-09-20 07:12 PT · full detail in the sections below — jump straight to one:

1. The pipeline in one line

2. The offering menu — four services, mix & match

3. Stage 1 — Lead generation

4. Stage 2 — Lead selection

5. Stage 3 — Cold outreach email

6. Stage 4 — Demo per lead

7. Stage 5 — Payment page → Stage 6 — Provisioning & operation

8. GAP ANALYSIS — Barry's described flow vs what exists (2026-09-19)

[Download PDF ↓](#)

DETAILS

End-to-end machine for AI Assistance / Apex Receptionist, stage by stage: what exists, where it lives, and where Barry's described flow diverges from reality (\$8 gap analysis).

Convention: discussion → update the plan. Same rule as `MASTER-PLAN.md` — after any discussion that changes this machine, update this doc in the same session.

Updated: **2026-09-19 ~10:15 PM PT** by Sage@miniPC **Standing locks:** SEND NO until Barry names the shop · Never Fuller · demo DID +1 (650) 476-2005 only · no purchases without Barry.

1. The pipeline in one line

Generate leads → select keeps → legal review → cold email (HELD → Barry gate) → prospect tries the demo → pays → we provision → we operate.

Cell chain that runs it: Leads → Legal → Marketing → QA → CoS hold → **Barry send/close** → Sage provision/operate.

2. The offering menu — four services, mix & match

#	OFFERING	WHAT THE CLIENT GETS	EXISTS TODAY?	PRICE
1	AI voice assistant	Their number answered after hours: caller greeted, qualified, booked to next open slot, window texted to the shop	Closest to real. Retell staging agent live on demo DID; 8 vertical flows; disclosures in greeting; but no per-customer number, and booking SMS/persistence are stubs (G4/G5)	\$199/mo decided (Bay \$249 · test-city \$99)
2	AI text assistant (SMS / WhatsApp)	Same assistant over text: missed-call text-back, SMS/WhatsApp Q&A + booking	Not built (G1). Demo shows owner-side SMS/WhatsApp stubs only. WhatsApp needs Business API approval; SMS needs Twilio + toll-free verification or A2P 10DLC (days-weeks, needs entity/EIN)	not priced
3	Web app	Booking/ops web app for the shop (book engine, owner board, morning recap)	Demo only (G2). <code>webapp/</code> + <code>iphone/</code> run 8 verticals on a shared book engine — browser-session state, no persistence/auth/multi-tenant	not priced
4	Website	A proper website for shops that have none/bad ones (often the wrapper the assistant lives in)	Planned only (G3). We build demo sites for fake shops — nothing exists as a client-website offering	not priced
	Bundles	Clients mix & match (e.g. voice + text; website	Bundle pricing not defined — Barry input	—

#	OFFERING	WHAT THE CLIENT GETS	EXISTS TODAY?	PRICE
		+ web app + voice)		

3. Stage 1 — Lead generation

- **Sources:** official business sites only — no scraped directories, no purchased lists. Agent curls the real site, reads page text/JSON-LD for printed hours, owner names, published email.
- **Criteria (ICP):** owner-run small shops in after-hours-sensitive trades (plumbing, HVAC, electrical, roofing, auto); printed owner name; printed hours with a real dark window; no 24/7 claim. Finding that repeats every metro: the yield is in **suburb family shops** — metro cores all sell 24/7.
- **Kill rules (any one = drop):** 24/7/emergency claim · no owner name printed · chain/mega-op · Jobber/HCP/ServiceTitan AI already answering · guessed email.
- **Batch process:** 3–6 looks per pass (Houston ran 16), new metros preferred, dossier per keep (shop, owner-as-printed, hours, dark window, contact, pain quote, track, QA stamp).
- **Where:** materials page `website/projects/ai-assistance/leads/materials/` · log `agents/leads/progress.md` · registry `agents/leads/leads-data.json` (drives the status page via `bin/update-leads-status.py`).
- **Current inventory (2026-09-19):** 33+ researched · **17 keeps / 5 batches** (Austin, Columbus/Denver, Jax/Indy/Louisville/Omaha, DFW, Houston) · 11 Track A · 6 Track B parked · 0 sent.

4. Stage 2 — Lead selection

- **Track A** = business email printed on the official site → eligible for the email pipeline.
- **Track B** = phone-only → parked inventory, never dialed/texted.
- Selection today is **rule-driven by agents** (kill rules + tracks), not scored, and Barry does not hand-pick keeps — he gates only the send (see gap G6). Near-misses are logged for possible resurrection (e.g. a near-miss Houston plumbing shop via license lookup).

5. Stage 3 — Cold outreach email

Template anatomy (see `dfw-track-a-emails.md`, `houston-track-a-emails.md`):

1. **Greeting** — owner's first name(s) exactly as printed on their site.
2. **Their own hours, quoted back** — "Your contact page says Monday to Friday, 8:00 to 5:00..."
3. **Missed-business math** — the 5:15 PM burst pipe hits voicemail until morning; we answer, book next open, text the shop the window. Guy-safe default: 5:15 PM / next-open / SMS.
4. **Price + payback** — "\$199 a month. One extra completed job pays the month."
5. **Online convenience** — demo hub link (Shop → Book → Card → Stubs) + staging voice line, both explicitly disclosed as a **fake shop, not their line**.
6. **CAN-SPAM footer** — commercial-message disclosure, postal address (`[POSTAL ADDRESS – Barry to insert]` — QA blocker #1),

unsubscribe honored ≤ 10 business days.

Legal conditions (per-shop review, `houston-track-a-legal-review.md` + DFW equivalent): verified role email re-checked live on review day · truthful subject · no misstating the prospect's own hours/services (DTPA — e.g. Lead D (Houston-metro plumbing) sells after-hours, so capture-layer angle only) · TCPA/DNC not triggered (email-only) · named do-not-dial numbers per shop.

HELD / authorization gate: drafts live on disk, CoS holds the queue, **AUTHORIZE SEND: NO** until Barry names the shop. Sends go from **Barry's Gmail**, one shop per email. Nothing has ever been sent. (A parallel gated tool exists at `receptionist/email-automation/` — dry-run only, triple-gated; see G8 for the divergence.)

6. Stage 4 — Demo per lead

- **Today:** one generic demo per **vertical** (8: plumbing, HVAC, construction, auto, restaurant, dental, laundry, law) at the hub + one shared staging voice line. Emails link the vertical closest to the shop.
- **Designed but not built:** per-lead custom demos. The email-automation template carries `{DEMO_VOICE_URL}/{DEMO_TEXT_URL}/{DEMO_WEB_URL}` with a `?business=NAME` idea, but the demo engine ignores such a parameter, and there is no clone-template-for-lead flow (G5). The `retell-flow` copy-template-per-client architecture is the intended base.

7. Stage 5 — Payment page → Stage 6 — Provisioning & operation

- **Payment:** the hub's "Stripe checkout path" is an explicit sketch — no Stripe account, key, card field, or charge (G4). Plan: Barry creates the account (Batch #1) → Sage wires a \$199/mo subscription Payment Link + signup webhook.
- **Provisioning (designed, waiting on Batch #1):** paid signup → buy client DID (Barry auths) → clone vertical template with their greeting/hours/escalation → disclosures on → client sets conditional forwarding from our one-pager (`onboarding-runbook.md` , `customer-intake-form.md`) → joint test call → 48 h burn-in.
- **Operation:** call-path monitor (`aiassist-callpath-watch` , live for the demo DID, 15-min cycle) extends per client · cost model `cost-model.md` (100 calls/mo ≈ \$42 cost → \$157 margin) · daily transcript review during burn-in · fallback forwarding (designed, per-customer config).

8. GAP ANALYSIS — Barry's described flow vs what exists (2026-09-19)

Ranked by damage to the sale.

#	GAP	BARRY'S PICTURE	REALITY IN THE REPO	FIX PATH
G4	Core promise is stubbed	"We answer, book, text you	Owner SMS is a stub; bookings live only in the browser session; no backend	Plan 1.2–1.3 after Barry Batch #1 (Stripe + Twilio + entity for A2P)

#	GAP	BARRY'S PICTURE	REALITY IN THE REPO	FIX PATH
		the window"	store; Stripe is a sketch (no account)	
G5	Per-lead custom demos don't exist	Each lead gets a custom demo link made from the vertical templates	Only generic vertical demos + shared staging DID. ? business=NAME personalization is in the email template's vocabulary but the demo engine ignores it; no per-lead cloning flow	S-M: query-param personalization (shop name/hours injected into demo) is a day; true per-lead agent clones ride the retell flow template (plan 1.4/3.2)
G1	Text assistant (SMS/WhatsApp) not built	A sellable offering #2	Demo stubs only. No outbound SMS (Twilio account missing, A2P 10DLC needs entity/EIN, days-weeks); WhatsApp Business API approval never applied for	Start registrations in Barry Batch #1; build after voice v1 (Phase 4)
G2	Web app offering = demo only	Sellable offering #3	8 vertical webapps on a shared book engine, but browser-session state only — no persistence, auth, multi-	Phase 4 (multi-tenant backend); do not sell before customer #1 exists on voice

#	GAP	BARRY'S PICTURE	REALITY IN THE REPO	FIX PATH
			tenancy, or prod deploy	
G3	Website offering doesn't exist	Sellable offering #4	No client-website product anywhere in the repo — only our own demo/hub pages prove the craft	Define scope + price with Barry; cheapest of the four to pilot manually
G6	Lead selection is rules, not Barry, and not auto	"Manual by Barry / auto rules"	Barry never picks leads — agents apply kill rules; no scoring model, no auto-picker. Criteria ARE documented (materials page + leads-data.json), so the gap is the picker, not the criteria	Keep rule-driven for now; add scoring only when volume demands it
G7	Pricing exists for 1 of 4 offerings	A menu clients mix & match	Only voice is priced (\$199/\$249/\$99). Text/web-app/website and bundles unpriced — outreach can't mention them	Barry decision (MASTER-PLAN "needs from Barry" #5)
G8	Two competing send mechanisms	One outreach machine	Hand-written per-shop drafts (current convention, Barry's Gmail)	Declare hand-drafts canonical for Track A keep email-automation mothballed or re-

#	GAP	BARRY'S PICTURE	REALITY IN THE REPO	FIX PATH
			AND email-automation/ bulk tool (dry-run, triple-gated, its own leads.json now stale vs leads-data.json)	point it at leads-data.json
G9	Draft inventory drift — RESOLVED 2026-09-19	"11+ drafts held"	Only the 4 DFW/Houston drafts exist on disk. Investigation (Sage, 2026-09-19): austin-track-a-emails.md never existed — not in apex git history on any branch, not on miniPC or VPS disk, no draft copy in hub HTML. The claim originated in the 2026-09-04 Grok-CoS handoff docs; any drafting happened only inside an unpersisted Grok session. Records corrected: the 4 early keeps (Lead E (Austin plumbing), Lead F (Austin	Plan 2.5: draft the 7 dossier-only Track A keeps via the now-standard Legal→Marketing→Q, chain

#	GAP	BARRY'S PICTURE	REALITY IN THE REPO	FIX PATH
			plumbing), Lead G (Columbus auto), Lead H (Columbus auto)) reset to dossier in <code>leads-data.json</code> ; page regenerated	
G10	No production call path (from readiness audit, still true)	Real clients get their own number	One staging agent + one shared demo DID; no per-customer DID purchase or agent provisioning; no fallback if Retell dies mid-call	Plan 1.4/3.2 + O2 fallback config

What already matches Barry's picture: the lead pipeline machinery (cells, tracks, kill rules, legal reviews, HELD gate) is real and disciplined; the demo hub is a genuine sales asset; the voice agent answers a real number today; compliance prep (disclosures, legal drafts, runbook, monitor, cost model) landed 2026-09-19.

Web-version note: specific lead business and owner names are redacted here to “Lead A-K (metro, vertical)” labels. Full names live in the private repo registry (`agents/leads/leads-data.json`) and the live **leads pipeline**.

then `rsync docs/` per `COS-DEPLOY.md` . Never Fuller.