EXPERT AI PROMPTS · VOICE DNA SYSTEM · TOOL 1

Operator's Field Instructions

Voice-First Client Onboarding Workflow System

How to run a client interview: what to check before you start, what stays out of the conversation, how to capture an answer accurately, and where the files go afterwards.

Field	Detail
Version	V2.2
Published	12 September 2026
Supersedes	V2.1
Applies to	expertaiprompts.info/voiceonboardingtool/ — Build 2.2.0
Read with	The Manual Capture Method (voice dictation guide) — the capture method in full. The Local Model Setup Guide — running the sensitive stages offline. Both are linked from the tool page.
Audience	A consultant capturing a client's material under an engagement. Self-serve users working on their own material need only the short form printed on the tool page itself.
Next review	March 2027, or sooner if an assistant provider changes its data-handling position.

What changed in V2.2

Session timings have been added. V2.1 twice warned about time pressure without ever stating how much time a session takes. Section 4 now carries a per-stage allowance, section 4.2 explains how it is made up, and the session checklist carries it as a column.

Stage 1 is named as the longest. It carries seven of the twenty-one mandated dimensions where the others carry six, five and three, and it runs close to twice Stage 4, so it is no longer safe to plan four equal sittings.

Step F has been given a budget. Reading and correcting accounts for roughly three minutes of every stage, and saying so is a better defence of the step than describing it as important.

Section 2.4 now states the scheduling consequence of running a stage on the local model: the stage has to be finished in one uninterrupted block, and how long that block is.

Applies to Build 2.2.0, which publishes the same timings in the tool page copy.

1. Before you start — three rules

Do these before the first stage prompt is copied. They take about a minute the first time and seconds thereafter, and they are the difference between a workflow that can be described to a client and one that cannot.

1.1 Rule 1 — Check the training setting. Paying for the account is not the check.

On consumer plans a paid subscription buys capacity, speed and features. It does not buy an exemption from model training — free and paid consumer accounts sit in broadly the same position on that question, and the control is a setting inside the account rather than the price of it.

Before a client session, open the account's data or privacy settings and confirm that training on your conversations is switched off. Where the work is done under a business, team or enterprise plan, customer content is generally excluded from training by default and there is nothing to switch.

Do this once per account and re-check it after any plan change, because these settings have been altered by provider policy updates more than once. If you cannot confirm it, treat the session as if Rule 2 applies to everything.

Plan tier	Position on training	What to do
Free consumer	Conversations may be used to improve models. An opt-out setting is generally available in the account.	Usable for Stages 1 and 4 with Rule 2 applied. Turn the setting off first.
Paid consumer — Plus, Pro, Max and equivalents	Broadly the same position as free. The subscription buys capacity, speed and features; it does not by itself buy an exemption.	Same as free. Do not treat the subscription as the control — check the setting.
Business, Team, Enterprise	Customer content is generally excluded from training by default, under the contract rather than a toggle.	This is the tier for client interview material. The exclusion holds without anyone remembering to set anything.
API and developer workspaces	Generally excluded from training by default.	Fine, but a poor fit for a voice interview workflow.
Locally hosted model	Nothing leaves the machine, so the question does not arise.	See section 2.4 and the Local Model Setup Guide.

Three things that follow from this table

A per-account setting is a weaker control than a contract. It can be reset by a policy change, it does not travel to a second device or a second account, and nothing on screen reminds you to re-check it. That is a reason to prefer a business or enterprise plan for client work, not a reason to distrust the setting.

"We use paid accounts" is not a claim to make to a client or an assessor. It sounds like a data-protection statement and is not one. Name the plan type and the setting, or say nothing.

Provider positions moved during 2025 and 2026 and may move again. Verify the current setting in the account rather than relying on a policy summary — including this one, which is a category description and not a statement about any named provider on any given date.

1.2 Rule 2 — Three things never go into a cloud assistant

These are the categories the wider architecture routes to a locally hosted model. The voice interview sits upstream of that routing, which is why the rule has to be applied by the operator here.

- **Named third parties.** Real prospects, customers and competitors, by name. Use a role and a sector instead: "a works manager at a regional council", not the person's name.
- **Financial figures.** Pricing, margins, contract values, revenue. Describe the shape without the numbers: "priced above the market on a value argument", not the rate.
- **Competitive strategy.** Anything that would damage the client if a competitor read it.

If any of these is essential to the artefact, capture it separately and add it locally afterwards — or run that stage on the local model under section 2.4. It does not belong in a cloud interview.

Say this to the client at the start, in one sentence, rather than interrupting them later. Most clients volunteer a real customer name within the first two minutes of Stage 2 if nobody has told them not to.

1.3 Rule 3 — Read the transcript before you send it

Use dictation mode, not live voice mode, so the text appears on screen before it is submitted. Read it back from the screen, type over any wrong name or number, and check that the answer actually addresses the question that was asked.

The full method — the custom word list, the Speak–Read–Fix–Send loop, the three-question answer check and the cleanup pass — is in the Manual Capture Method, and it is not optional. Every artefact produced here becomes an input to every prompt that follows, so an error accepted at this stage travels through everything downstream and becomes progressively harder to see.

2. Your capture method and platform

This section replaces the Google Gemini voice-app interface guide in V2.0. The method is no longer described in terms of one vendor's buttons, because the method does not depend on one vendor.

2.1 Dictation mode, not live voice mode

Live voice mode and dictation mode are not two routes to the same place. In live voice mode your speech is transcribed and submitted in one motion — there is no interval in which a transcript exists on screen and has not yet been sent, which means there is no point at which a name or a figure can be corrected before the model sees it. The review step in Rule 3 is therefore not merely skipped in live voice mode. It is unavailable.

Dictation mode places the transcribed text in the message box and waits for you to press send. That pause is the control this entire system is built on.

> **The one exception, and its consequence**
>
> Live voice mode is acceptable for exploratory or self-directed use where no client material is involved and no artefact is being produced for a third party — talking through Stage 1 about your own background, for example.
>
> If it is used at any point in a client session, the transcript must be read back and corrected in the dashboard textarea before Save as Text File. That preserves the control by moving it one step later. What is not acceptable is using live voice mode and skipping the correction entirely.

2.2 Which assistant

Three assistants are supported: Claude, ChatGPT and Google Gemini. These are the three the Manual Capture Method carries setup instructions for, and that is the selection rule — an assistant without published dictation instructions is not a supported assistant, because the operator has no way to confirm they are in the right mode.

Do not substitute a fourth on your own initiative during a client session. If one should be added, it is added to the dictation guide first.

2.3 Recommended platform — the Claude desktop app

The Claude desktop app is the documented default for client work, for two capability reasons rather than a preference.

- **Dictation is built in.** The transcript appears on screen without you selecting a mode, which makes the Rule 3 requirement structural rather than something you have to remember. A control the interface enforces by sequence is stronger than one that depends on the operator choosing correctly under time pressure in front of a client.
- **It can write files.** Exports can land in the client's folder rather than in Downloads, which closes the file-destination problem in section 5.3 at source rather than by instruction.

This is a platform recommendation, not a tier position. A desktop app on a personal plan sits in exactly the same training position as the browser — Rule 1 still applies, and a business or enterprise plan is still the answer for client work.

2.4 When the material cannot be abstracted — the local model

Stages 2 and 3 are the two that elicit named prospects, pricing and competitive strategy. Where Rule 2 cannot be satisfied because the artefact genuinely needs the specifics, run those stages on the locally hosted model instead: speech to text on your own machine, then the corrected transcript into a model that never sends anything anywhere. Setup is in the Local Model Setup Guide.

Stages 1 and 4 carry far less sensitivity and are the stages where a frontier model's output quality is most worth having, so a split is legitimate and is usually the better choice.

> **Decide before the stage starts, and finish where you started**
>
> A stage that runs locally by default but falls back to a cloud assistant because the local model is slow or the answer is unsatisfying gives you the exposure and the paperwork saying you do not have it — which is worse than running the stage in the cloud and documenting it honestly.
>
> The local path for Stages 2 and 3 is recommended rather than mandatory. An output-quality comparison against the cloud version is planned and this section will be updated with what it finds. Allow for the scheduling consequence as well: the local model keeps no conversation history, so a stage run locally has to be finished and copied out in one sitting. Stage 2 is about twenty-one minutes and Stage 3 about seventeen, so running either locally means protecting an uninterrupted block of that length — harder to guarantee in a client's office than at your own desk.

3. The workspace

V2.0 required a two-window arrangement: the companion dashboard on one side of the screen, the assistant on the other. That remains a good arrangement and is still the right answer in a browser, where you are moving text between two tabs and want both visible.

It is no longer mandatory. Under the platform recommendation in section 2.3 the assistant is a desktop application rather than a second browser tab, and a two-window layout becomes a preference rather than a requirement. Arrange the screen however lets you read a transcript comfortably and paste without losing your place.

What has not changed is the reason the arrangement existed: you must be able to see the transcript before you send it, and see it again in the dashboard textarea before you save. Any layout that makes either of those awkward is the wrong layout.

4. The 4-stage context chaining pipeline

The stages run in order and each one feeds the next. Do not run them out of sequence, and do not skip one because the client is short of time — a missing stage does not simply remove its own artefact, it degrades every artefact after it. Allow about seventy minutes for the full set, and see the per-stage allowance in the table below: they are not equal, and Stage 1 is nearly twice Stage 4.

Stage	What it captures	Framework	Sensitivity	Allow
1 — Who Am I	Personal Value Proposition: background, motivation, working style, what you are actually good at.	Bolles 7-dimension Parachute framework	Low	22 min
2 — Customer	Customer Persona: the buyer's role, responsibilities, pains, selection criteria, yes-triggers and objections.	Buyer psychology deep dive	HIGH — elicits real prospect names	21 min
3 — Value Prop	Business Value Proposition: measurable results, referral evidence, moats.	5 Core Proof Questions	HIGH — elicits pricing and competitive position	17 min
4 — Brand Identity	Verbal tone, visual guardrails, what the brand will never sound like.	3 Voice and Brand Questions	Low	12 min

> **Stages 2 and 3 are named deliberately**
>
> Rule 2 applies to all four stages, but these are the two where a client will volunteer a category-one answer without being asked for one. Stage 2 asks them to describe their buyer in detail, which surfaces real prospects by name. Stage 3 asks about moats and evidence, which surfaces pricing and competitive strategy.
>
> Naming the stages rather than saying "sensitive material" is deliberate. A rule that names its scope is auditable; a rule that does not is a matter of opinion at the moment it is applied. This is the same convention used for the four mandated-local workflows in the CDD chaining workflows set.

4.1 Running a stage — Steps A to G

Seven steps. Step F is the one that most often gets skipped under time pressure, which is exactly why it is a numbered step rather than a note. It is also budgeted: reading and correcting the stage document is about three minutes of each stage, and checking each answer before you send it accounts for a further two to three. Across the four stages that is thirty of the session's seventy minutes. It is not overhead on the work; it is close to half of it.

Step A Set the client details

Enter Your Name / client person, Company or Brand Name, and confirm the Date Run at the top of the dashboard.

Do this first. These fields name every export, and a file saved before they are set is hard to tell apart from another client's later.

Step B Copy the stage prompt

Open the stage accordion and click Copy Prompt. Do not retype or edit the prompt — it is written to produce a specific structured output that downstream tools parse.

Step C Run the interview in your assistant

Paste the prompt into the assistant and answer the coach's questions in your own words, one at a time.

Use dictation mode. Paste the client word list first if you have one — proper nouns and figures are where transcription fails.

Apply Rule 2 as you go. Describe the shape, not the specifics.

Step D Copy the synthesised output

When the interview completes, copy the assistant's final document: the narrative, the structured JSON block, and the outstanding-questions backlog.

Take the whole thing. The dashboard parses the JSON block out of it and the downstream assemblers use both halves.

Step E Paste it into the stage textarea

Paste into the transcript box for that stage on the dashboard.

Step F Read it and correct it

Read the pasted text before you save. Fix any wrong name, figure or term directly in the textarea.

Check the answers against the questions — a fluent answer to a question nobody asked is the failure mode that survives longest, because it reads well.

The character-count indicator underneath is a depth guideline only. It measures how much was said, not whether it is right, and it never blocks saving. It is not this step.

Step G Save as Text File

Click Save as Text File. The next stage opens automatically.

The file downloads to your usual Downloads folder — see section 5.3 for where it goes next.

4.2 How long to allow

About seventy minutes for the full set of four, of which roughly thirty is quality control rather than conversation. The per-stage allowance is in the table in section 4, and the four are not equal.

Each stage divides the same way. Interview time is reading the coach’s question, thinking, and dictating the answer — about ten minutes a stage on average. QA time is checking each dictated answer before it is sent, reading and correcting the stage document afterwards, and the copy, paste and save around it — about seven or eight minutes a stage.

The tool page states the same model to a self-serve user as about ten minutes per interview, plus about thirty minutes across the set to check. That is the same seventy minutes counted the other way. This document gives totals per stage because an operator books wall-clock.

Stage 1 is the longest at about twenty-two minutes. It carries seven of the twenty-one mandated dimensions where the others carry six, five and three, it generates the most follow-up questions, and its framework asks for real-world stories in plain trade language, which run long. Book it on its own. Stages 3 and 4 together come to about twenty-nine minutes and sit comfortably in one sitting; Stage 1 does not leave room for a companion.

These figures assume dictation, which is the documented capture method. They are modelled from the four stage prompts rather than measured from a timed run, so treat them as planning figures and revise them once a real session has been clocked.

5. Operational data management

5.1 Everything stays in your browser

The dashboard uses HTML5 local storage. Nothing is transmitted to a server and there is no database, so a browser crash or an accidentally closed tab does not lose the session — reopen the page and the transcripts are still there. This is a materially better privacy position than most comparable tools and it is worth stating to a client in those terms.

The other half of that fact matters just as much: local storage is unencrypted and readable by anything with access to that browser profile. It is a good place for a working session and a poor place for anything to sit for a week.

5.2 Clear All Data — before and after

The red Clear All Data button at the top of the dashboard wipes the stored session: both name fields, the date, and all four transcripts.

- **Before a new client.** This is what V2.0 documented, and it is still right. It protects the previous client from appearing in the next session's exports.
- **After your own session.** New in V2.1, and the more important of the two. Clearing before a session protects the previous client; clearing after protects the current one. The material is legible to whoever next opens that browser profile.

Do both. Save every stage to a text file first — clearing is not reversible.

5.3 Where the files go

Exports auto-name and download to your operating system's Downloads folder. Move them into the client's own folder as soon as the session ends. They contain the client's value proposition, evidence and buyer detail in plain text and unencrypted, and Downloads is a shared, rarely-reviewed location on most machines.

Under the platform recommendation in section 2.3 the assistant can write directly into the client folder, which removes this step rather than requiring you to remember it.

5.4 Retention

Interview exports and the artefacts derived from them are retained for the life of the engagement plus seven years, then destroyed.

State this if a client asks. An unstated retention period is the one that causes trouble, because it invites the assumption that files are kept indefinitely and everywhere.

6. Integrating with downstream prompt assemblers

Unchanged from V2.0. The export contract has not moved and every downstream tool continues to parse these files without modification.

Item	Detail
Filename convention	[Company]_[Person]_[DocumentType]_[DateRun].txt

Item	Detail
File structure	A human-readable record — the interview prompt used, the raw transcript, and the synthesised output — followed by a machine data block containing the parsed JSON.
Finding the JSON	Look for the machine data block fence near the end of the file. Everything between the fences is the structured data.
Where it goes	Paste the JSON block into the Expert AI Prompt Assembler or the CDD Blueprint Builder, or hand the whole file to a marketing agency. The full file is the more useful hand-off, because the narrative carries nuance the JSON does not.
Do not edit	The machine data block. Correct errors in the transcript before saving (Step F) rather than editing the JSON afterwards — the two must agree.

7. Session checklist

One page to work from once you know the material.

When	Do	Allow
Before the session	Confirm the assistant's training setting is off, or that you are on a business/enterprise plan. Decide which stages run locally, if any. Prepare the client word list. Click Clear All Data.	5 min
Opening the session	Tell the client in one sentence what stays out of the interview: real names, real numbers, competitively sensitive strategy.	1 min
Each stage	Set client details (A). Copy the prompt (B). Run the interview in dictation mode (C). Copy the whole output (D). Paste it in (E). Read and correct it (F). Save as Text File (G).	22 / 21 / 17 / 12 min
Each answer	Read the transcript on screen. Fix names and figures. Check the answer addresses the question. Then send.	about 30 sec each
After the session	Move exported files from Downloads into the client folder. Click Clear All Data. Confirm all four stages saved.	5 min
Later	Retain for the engagement plus seven years, then destroy.	—

8. Version map and status

Sections were renumbered in V2.1 to make room for the pre-flight rules. This is the map. The last four rows record what V2.2 changed; those sections keep their V2.1 numbers.

V2.0 section	V2.1 section	Change
—	1. Before you start — three rules	New. Training setting and tier table, the three categories, read before sending.
1. Google Gemini Voice App Interface Guide	2. Your capture method and platform	Rewritten. Dictation mode replaces the Gemini voice-app walkthrough; assistant list, platform recommendation and local-model option added.
2. The Dual-App Workspace Protocol	3. The workspace	Demoted from mandatory to preference.
3. The 4-Stage Context Chaining Pipeline	4. The 4-stage context chaining pipeline	Stages 2 and 3 named as sensitive. Steps A–F become A–G, with a read-and-correct step inserted at F.
4. Operational Data Management	5. Operational data management	After-session clearing, file destination and retention added.
5. Integrating with Downstream Prompt Assemblers	6. Integrating with downstream prompt assemblers	Unchanged in substance.

V2.0 section	V2.1 section	Change
—	7. Session checklist	New.
—	8. Version map and status	New.
—	2.4 When the material cannot be abstracted	Scheduling consequence of the local model added: how long the uninterrupted block is.
—	4. The 4-stage context chaining pipeline	Per-stage allowance added to the stage table. New 4.2 How long to allow.
—	4.1 Running a stage — Steps A to G	Step F given a time budget.
—	7. Session checklist	Allow column added.

Status

This document describes how to operate a documented workflow. It is not legal advice and not insurance advice, and where it refers to data-handling concepts it does so to explain the operating rule rather than to state what any obligation requires.

The plan-tier descriptions in section 1.1 are category distinctions, not statements about any named provider on any given date. Provider positions moved during 2025 and 2026 and may move again — verify the setting in the account rather than relying on this summary.

The local-model recommendation in section 2.4 is based on which stages elicit sensitive material, not on a completed output-quality comparison. That test is planned and this document will be updated with its result.