

Sample Probe Report

Four probes of forty-three, shown in the exact format every client report uses.

Where these findings come from. Every probe below was run against our own production receptionist, the same live line printed at the bottom of this page. Each one found a real defect, and each was fixed before we sold this audit to anyone. Client reports carry per-probe evidence exactly like this: what we did, what happened, the fail signature, and the fix, ranked by severity.

S1 a real caller loses something: a booking, their data, the call **S2** the agent degrades in a way the caller notices

S3 the operator is blind to something they need to see

A5 TRANSPORT AND TOOL WIRING

S1

PROBE · Diff the tools attached to the live agent against the intended set, then watch the endpoint for calls with no matching handler.

WHAT WE FOUND · A renamed capture tool was retired from the codebase but never detached from the agent. The model kept choosing the stale tool by its cleaner name. Its calls arrived as bare, unauthenticated requests, and the real capture path never fired. The line looked healthy for days while it booked nothing.

FAIL SIGNATURE · Endpoint receives calls it has no handler for; the destination table stops growing while every dashboard stays green.

FIXED · attached-tool diff now runs on every release. When you replace a tool, you kill the tool.

A7 TRANSPORT AND TOOL WIRING

S1

PROBE · Call every tool endpoint from a public vantage point, never from inside the operator's own network.

WHAT WE FOUND · The public funnel was down for four days while every internal health check reported green. The agent answered, carried the conversation, and tried to book; every tool call 500ed at the public edge. All monitoring ran from inside the network, so nothing could see what a real caller hit.

FAIL SIGNATURE · Every tool 500s at the edge; zero entry logs; internal checks green the whole time.

FIXED · external synthetic probe on a schedule. A check that shares the system's vantage cannot see this class.

D6 CONVERSATION CONTROL

S2

PROBE · Read the platform's latency-filler configuration and the prompt's bridging instructions side by side, then place a multi-tool call and count time-buying phrases per turn.

WHAT WE FOUND · The filler had two owners. The platform shipped a randomized stall-phrase pool with a 1.5 second trigger, and the prompt separately mandated an acknowledgment on every turn. Neither owner could see the other. One live turn produced five stall phrases before any content, and a filler prefixing the goodbye.

FAIL SIGNATURE · Multiple time-buying phrases in one turn; a filler before a reply that needs no lookup at all.

FIXED · one owner. The prompt narrates the action, one line per turn; the platform pool is demoted to a single long-fuse backstop. Verified on a live pass transcript: zero stalls.

I5 OBSERVABILITY AND EVIDENCE

S2

PROBE · Write a config change through the platform API, then read the live object back field by field and compare against what was sent.

WHAT WE FOUND · The console accepted a nested partial update, returned 200, and silently dropped the field. A per-generation cap was sent three ways and acknowledged every time; the agent kept running the old value until the write carried the full parent object.

FAIL SIGNATURE · Post-write read-back differs from what was sent, with no error anywhere.

FIXED · every config write is followed by a field-by-field read-back comparison. The 200 is not the receipt; the read-back is.

The nine failure classes the full audit covers: A transport and tool wiring · B identity, auth, and secret hygiene · C idempotency and state races · D conversation control · E truthfulness and commitment boundaries · F adversarial input · G failure and degradation · H data protection and compliance · I observability and evidence. Forty-three probes, twenty-six automated, each with a fail signature and the usual fix.

The full audit: \$5,500 fixed, five business days. Graded report with per-probe evidence, a ranked fix list, and a regression checklist your team keeps. Unconfirmed findings are dropped, not hedged. We run our own receptionist on this discipline: call **+1 (617) 766-0577** and try to break it. · gnosislabs.tech/voice-audit · Cap Dawes, Gnosis Labs