

PUBLIC SAMPLE FORMAT

Monthly Health Report

A post-launch operating report for an AI agent fleet: usage, errors, drift, costs, tool health, completed tune-ups, open risks, and the next recommended fixes.

█ SAMPLE ONLY - NOT A CLIENT REPORT



AUDIENCE

Retainer client and operator owner

PURPOSE

Keep launched systems healthy

OUTPUT

Metrics, fixes, risks, next actions

WHAT WE WOULD REVIEW EACH MONTH

Agent fleet health

This sample assumes a generic post-launch client with one or more AI agents running inside real workflows. The real report would use the client's logs, tickets, alerts, usage, and approved business context.

Health rule: a launched agent is not finished forever. It needs monitoring for tool breakage, model changes, prompt drift, bad edge cases, cost spikes, and human review feedback.

USAGE

Is the system being used?

The report checks runs, skipped runs, review volume, completion status, and whether the agent is helping the workflow it was built for.

QUALITY

Are outputs still trustworthy?

Review notes, corrections, failed cases, missing context, and hallucination flags show where the agent needs tuning.

RELIABILITY

Are tools still connected?

The report checks API failures, auth issues, webhook changes, schema changes, broken forms, and stuck queues.

COST

Are model costs drifting?

Usage and token costs should be monitored so an agent does not quietly become too expensive for the value it creates.

Health fields we confirm

LOGS

Agent runs, tool calls, source links, human review notes, errors, skipped jobs, and edge cases.

COSTS

Model usage, tool usage, hosting, monitoring overhead, and any vendor changes that affect monthly cost.

NEXT FIXES

The tune-ups, guardrails, prompt changes, integrations, or workflow changes recommended next.

WHERE THE FLEET STANDS

Monthly operating status

The report translates logs into plain English: what ran, what failed, what changed, what was fixed, and what needs a decision from the client.

STATUS POSTURE

Keep the fleet narrow and observable

The monthly report should make it easy to decide whether to tune the current workflow, expand the agent, pause an unsafe path, or build the next narrow system.

HEALTHY SIGNALS

What good looks like

- Runs complete as expected
- Outputs show sources and review status
- Human corrections trend into tune-ups
- Costs match usage and value

UNHEALTHY SIGNALS

What needs attention

- Repeated missing context
- Tool errors or auth failures
- Unexpected model cost spikes
- Reviewers no longer trust outputs

WHAT WOULD GO INTO THE REPORT

Monthly report components

The report should focus on what changed, what was fixed, what remains risky, and what decision should happen next.

01 **RUN SUMMARY.** Show how many times the agent ran, what jobs completed, what skipped, and which workflows used it.

02 **ERROR REVIEW.** Summarize tool failures, bad inputs, missing records, review corrections, and recurring edge cases.

03 **COST REVIEW.** Note model usage, tool cost, hosting cost, and whether the spend still matches the workflow value.

04 **TUNE-UPS COMPLETED.** List prompt changes, tool fixes, guardrail updates, workflow edits, and test cases added.

05 **NEXT ACTIONS.** Recommend what to keep, fix, pause, expand, or build next based on the evidence.

The point: post-launch support should make the system more useful over time instead of leaving the client with a black box.

OPERATING RISKS

Risk flags we would call out

A useful health report is honest about places where the system is drifting, breaking, costing too much, or depending on fragile human habits.

TOOL DRIFT

Integrations can change

Vendor UI changes, auth expiry, webhook updates, schema changes, and API limits can break a workflow that used to run cleanly.

PROMPT DRIFT

Outputs can get stale

Policies, offers, products, pricing, and team preferences change. Prompts and source rules need to stay current.

REVIEW FATIGUE

Human feedback may disappear

If reviewers stop marking corrections, the system loses the signal it needs for useful tune-ups.

SCOPE CREEP

The fleet can grow too fast

Adding new agents before the first ones are stable can create more monitoring work than value.

What the client gets next

BRIEF

A monthly readout covering runs, errors, drift, costs, completed tune-ups, open risks, and recommended next actions.

FIXES

Approved tune-ups move into the next retainer work cycle, with scope matched to the active support tier.

ROADMAP

If the current workflow is healthy, the report can recommend the next agent, integration, or automation improvement.