

Full-team audit — sample report.

What you receive at the readout, shown on a worked example:
a composite 34-person independent insurance agency.

Cerebrum One · cerebrumone.ai/sample-audit-report · runbetter@cerebrumone.ai

ILLUSTRATIVE SAMPLE

This is not a real client. "Meridian Coverage Group" is a composite built from published industry benchmarks and our audit methodology, so you can see the exact shape and rigor of the deliverable before paying for one. Every figure below follows the same math a real report uses — and in a real report, every figure is yours.

The engagement at a glance

Company	Meridian Coverage Group (composite)
Business	Independent insurance agency · commercial + personal lines
Size	34 employees — 12 service, 8 producers, 14 ops & admin
AMS	Applied Epic
Interviewed	31 of 34 employees (91%) · average 22 minutes each
Workflows described	61
Automation candidates	19
Cleared the materiality bar	7 (listed below)

Executive summary

We recommend building three automations. Together they recover roughly **36 service-team hours a week** — about **\$56,000 a year** at your fully loaded rates — for **\$35,000** in fixed-price builds: a blended payback of **7.5 months**. Of the remaining four candidates, one is a buy (a feature in software you already pay for), one is a process fix with zero build cost, and two are waits with revisit dates attached.

Total addressable waste identified across all seven: roughly **\$77,000 a year** — 8.1× the audit fee, and 2.7× the \$28,500 guarantee threshold this audit carried.

How the numbers are built

Challenge any input and the ranking recalculates. The assumptions:

- **Fully loaded hourly cost:** \$32/hour for service and admin roles ($\$24 \text{ wage} \times 1.33$ for taxes and benefits); \$38/hour for bookkeeping. In a real report these come from your payroll, not our defaults.
- **48 working weeks** per year.
- **Volumes** as reported in interviews, cross-checked against AMS activity logs and sent-mail counts where available. Where the interview and the system disagreed, we used the lower number.
- **"Automatable %"** is scored per workflow, not assumed. Each recommended build's basis is shown in its section below.

The ranked list

ALL SEVEN OPPORTUNITIES THAT CLEARED THE MATERIALITY BAR

#	OPPORTUNITY	HOURS/WK	LABOR COST/YR	AUTOMATABLE	CALL	FIXED QUOTE	PAYBACK
1	Certificate of insurance (COI) issuance	17	\$26,112	80%	Build	\$12,000	~7 mo
2	Renewal prep & remarketing	24	\$36,864	60%	Build	\$15,000	~8 mo
3	Carrier-portal re-keying into the AMS	10	\$15,360	85%	Build	\$8,000	~7 mo
4	New-business quote intake (ACORD re-key)	8	\$12,288	75%	Wait	est. \$9,000	~12 mo
5	Commission reconciliation	4	\$7,296	70%	Buy	—	—
6	Claims status follow-ups	5	\$7,680	50%	Wait	—	—
7	Client welcome packets	3	\$4,608	65%	Process fix	—	—

#5 is a native Applied Epic feature plus an \$89/month reconciliation tool — no build justified. #6 involves too much per-case judgment to automate safely today; revisit in two quarters. #7 is a document template and a mail merge in tools the agency already owns. #4 shares an extraction layer with #3 — see "What happens next."

Build #1 — COI issuance

Nine of the twelve service employees independently described the same workflow: a certificate request arrives by email, a CSR opens the policy in Epic, re-types holder details into the certificate, generates the PDF, and emails it back. About 85 requests a week, roughly 12 minutes each.

The math: $85 \text{ requests/wk} \times 12 \text{ min} = 17 \text{ hrs/wk} \times \$32/\text{hr} \times 48 \text{ wks} = \$26,112/\text{yr}$.

Scored 80% automatable (standard holders and wording are mechanical; nonstandard requests are not) → **\$20,890/yr recovered**. Fixed quote \$12,000 → **payback 6.9 months**.

How it works: the system reads the request email, matches it to the policy record in Epic, generates the certificate, and sends it. Anything nonstandard — holder wording changes, additional-insured requests — routes to a CSR for a one-click review. Nothing leaves the agency without either matching a known pattern or a human approval.

What stays human: new holders with custom wording, and anything that touches E&O judgment.

Build #2 — Renewal prep & remarketing

Six CSRs each spend about four hours a week assembling renewals: pulling the expiring policy, building the comparison, drafting remarket submissions.

The math: $6 \text{ CSRs} \times 4 \text{ hrs/wk} = 24 \text{ hrs/wk} \times \$32/\text{hr} \times 48 \text{ wks} = \$36,864/\text{yr}$. Scored 60% automatable (assembly is mechanical; coverage judgment is not) → **\$22,118/yr recovered**. Fixed quote \$15,000 → **payback 8.1 months**.

How it works: at 90 days out, the system pulls the expiring policy, builds the renewal comparison, drafts the remarket submissions, and queues the package for CSR review — so the CSR starts from a finished draft instead of a blank screen.

What stays human: every coverage recommendation and every client conversation.

Build #3 — Carrier-portal re-keying

Policy documents arrive in a dozen carrier portals and inboxes; someone re-keys the changes into Epic. Interviews put it at about 10 hours a week, spread across the service team — the "first hour of every morning" problem.

The math: $10 \text{ hrs/wk} \times \$32/\text{hr} \times 48 \text{ wks} = \$15,360/\text{yr}$. Scored 85% automatable (extraction and matching are mechanical; mismatches are not) → **\$13,056/yr recovered**. Fixed quote \$8,000 → **payback 7.4 months**.

How it works: the system watches the portals and inboxes, extracts policy changes, updates the matching record in Epic, and flags anything that doesn't reconcile for a human to resolve.

What stays human: every mismatch. The system never guesses at a discrepancy.

What we heard

Composite quotes, representative of what full-team interviews surface:

- "The first hour of my day is finding out what changed overnight — portal by portal." — CSR
- "I keep my own spreadsheet because I don't trust the dashboard. Everyone does." — Account manager
- "Renewals aren't hard. They're just the same eleven steps, ninety times a month." — Senior CSR

What happens next

The recommended sequence, if all three builds are approved:

- **Weeks 1–6:** Build #3 first. It is the fastest, and its document-extraction layer is shared plumbing for #1 and #4.
- **Weeks 3–10:** Build #1, overlapping once #3's extraction layer is standing.
- **Weeks 8–14:** Build #2.
- **Next quarter:** revisit #4 — with #3's infrastructure in place, its quote drops below \$9,000 and the payback math changes.

The \$9,500 full-team audit fee is credited in full against Build #1's invoice.

A real report looks exactly like this — with your numbers.

Full-team audit: \$9,500 flat, two to three weeks, refund guarantee in writing, fee credited toward your first build.