

AI Assessment: Findings Report

Client: Meridian Mechanical Ltd., Kelowna, BC

Prepared by Nathan Flavel, nAItē. Assessment dates: March 2 to 13.

This is a demonstration report for a fictional company. The method is exactly what a real client receives. Meridian Mechanical Ltd. is invented for this sample: a BC HVAC and plumbing contractor with 28 staff and roughly \$6.2M in annual revenue. No real engagement is implied.

1. What we did

We interviewed six people: the owner, the dispatcher and office manager (Dana), the bookkeeper (Priya), two journeymen (one plumbing, one refrigeration), and one apprentice. We sat in the office for two half-days with a stopwatch and timed the work as it happened. Twelve months of counts came from the quoting software, the accounting file, and the answering-service log.

Every number in this report comes from a count, a stopwatch, or your payroll.

Loaded-hour rates, built from Meridian's payroll:

- Office: $\$40,800 \text{ average wage} \times 1.30 \text{ on-costs (CPP, EI, WorkSafeBC, benefits)} = \$53,040 \div 1,710 \text{ productive hours} = \text{\$31/hr}$
- Journeyman: $\$71,000 \text{ average wage} \times 1.30 = \$92,300 \div 1,710 = \text{\$54/hr}$

The 1,710 figure is 2,080 paid hours minus vacation, stat holidays, sick days, and training time.

Where this report is weakest. Before the numbers, the two we trust least. First, the 62% quote follow-up share came from a four-week sample, not the full year. Second, the after-hours loss figure assumes the callers we never reached carried an average ticket and would have booked. Both are flagged where they appear. We also did not time permit paperwork ourselves. Dana's estimate is 35 minutes each. We watched one and it ran 31, so we kept her number.

2. The tally

Process	Volume (source)	Time each	Hours/yr	Rate	Annual cost
Quote follow-up	831 quotes chased (62% of 1,340 issued, share from the four-week sample), 2 touches avg = 1,662 touches	9 min (timed, range 6 to 13)	249	\$31	\$7,719
Invoice chasing	558 invoices past 30 days (18% of 3,100), 3 touches avg = 1,674 touches	7 min (timed)	195	\$31	\$6,045
Parts and PO phone	11 calls/day x 240	6 min (timed)	264	\$54	\$14,256

loop, field side	working days = 2,640 calls (one-week count, extrapolated)				
Permit paperwork	310 permits (count)	35 min (Dana's estimate, spot-checked once)	181	\$31	\$5,611
Timesheet keying and chasing	19 field staff × 52 weeks = 988 sheets	12 min (timed)	198	\$31	\$6,138
Re-keying completed jobs into QuickBooks	3,100 jobs (count)	6 min (timed)	310	\$31	\$9,610
After-hours enquiry loss	388 after-hours calls logged, 121 new-service enquiries, 44 never reached on callback (counts)	n/a	n/a	n/a	~\$9,900 margin (estimate: 44 × \$410 avg first-visit ticket × 55% gross margin)

Labour subtotal: 1,397 hours/yr, \$49,379. Total exposure including the after-hours estimate: roughly \$59,300 per year.

The after-hours row is the softest number on this page. Everything above it is a count multiplied by a stopwatch.

3. Addressable value

We never claim the whole row. A human still reviews each touch and handles the exceptions. Here is what remains after the haircut, and the basis for each cut.

Process	Annual cost	Basis for haircut	Addressable
Quote follow-up	\$7,719	Drafting automated, Dana still reviews each touch: 9 min drops to 2, saving 194 hrs	\$6,014
Invoice chasing	\$6,045	Same shape: 7 min drops to 2, saving 139 hrs	\$4,309
Parts/PO phone loop	\$14,256	We logged call reasons for one week: 61% asked for information already sitting in the office system. Only that share is addressable: 161 hrs	\$8,694
Permit paperwork	\$5,611	Forms prefill, but review and portal submission stay human: 35 min drops to 15, saving 103 hrs	\$3,193
Timesheets	\$6,138	Field entry removes keying, correction-chasing stays: 12 min drops to 3, saving 148 hrs	\$4,588
QuickBooks re-keying	\$9,610	Integration carries the data, Priya spot-checks: 6 min drops to 1, saving 258 hrs	\$7,998
After-hours loss	~\$9,900	We assume capture of half. Some of those 44 callers will not book with any system	\$4,950

Addressable total: roughly \$39,700 of a \$59,300 problem. 1,003 labour hours plus the after-hours capture.

4. The honest split

Of the \$39,700, only the after-hours capture, about **\$4,950 a year, is cash that stops leaving the business**. The other \$34,800 is 1,003 recovered hours. That is not profit. It is capacity, and capacity is only worth money if it is redeployed. At Meridian it has somewhere real to go: the 161 field hours flow into a job board that currently runs a two-to-three-week backlog, and the office hours go to the maintenance-contract renewals that lapsed last year because nobody had time to chase them (31 lapsed, by Priya's count).

5. What to build, in order

First, a free fix, before anyone pays anything. Give field staff read-only phone access to job records. That attacks the 61% of the parts/PO loop at zero cost. We will recount call volume after 30 days. We are deliberately not proposing a build for that row until the free fix has been measured.

Build 1: After-hours enquiry capture. Structured intake by SMS and web form that books a next-morning callback slot into the dispatch queue. It quotes no prices and promises no schedule. Cheapest build, and the only one that produces cash rather than capacity, so it goes first and earns trust for the rest.

Checkpoint: every enquiry lands in Dana's 7:30am queue for human callback.

Cost: \$2,500 build, \$95/mo. Expected return: ~\$4,950/yr.

Kill criteria: fewer than 8 booked callbacks in the first 90 days (baseline says 11 enquiries per quarter are being lost). If we miss that, we stop and remove this build's share from the fee.

Build 2: Quote follow-up and invoice-chasing drafter. One system drafts every follow-up and collection touch. Nothing sends itself.

Checkpoint: Dana approves and sends every message for the first 60 days, no auto-send.

Cost: \$4,500 build, \$140/mo. Expected return: ~\$10,300/yr in recovered office hours.

Kill criteria: at day 60, if median approval time exceeds 3 minutes per message, or fewer than 70% of drafts go out unedited, we stop and remove its share from the fee.

Build 3: Field timesheets with QuickBooks writeback. Field entry on the phone, data flows to payroll and job costing without re-keying. Most disruptive to daily habits, so it goes last.

Checkpoint: Priya spot-checks a 10% sample every pay run for the first three cycles.

Cost: \$6,000 build, \$180/mo. Expected return: ~\$12,600/yr in recovered office hours.

Kill criteria: if payroll corrections exceed the current baseline of 14 per month after one full cycle, or field submission is below 90% by week 6, we stop and remove its share from the fee.

6. What we recommend against

An AI phone receptionist on the daytime line. Do not buy one. Meridian's daytime line takes about 60 calls a day and most are existing customers who ask for Dana by name. She knows which furnace is on its last season and which customer needs the invoice explained gently. She is the customer relationship, and an AI answering that line would degrade the single strongest asset in the company to save money the tally shows is being lost elsewhere. The real after-hours gap is solved by Build 1 for \$95 a month.

We also looked at a website chatbot. The site log shows 9 chat-shaped enquiries a week. That volume does not support the build or the upkeep. The web form in Build 1 covers it.

7. Risk and governance snapshot

Four of the six people we interviewed are already pasting customer information into free AI accounts, including one instance of a customer's full service history with address. Nobody did anything wrong. There was no policy to follow.

Month one installs the three-pile data sort. Green: public information, any tool. Yellow: internal operations, paid accounts with training disabled. Red: customer PII, banking, employee records, which never enter an AI tool without a signed data agreement. It ships as a one-page policy, staff sign it, and the free accounts get closed the same week.

8. The decision page

Option 1: Do nothing. The tally says the current path costs roughly \$59,300 a year, about \$39,700 of it addressable. Doing nothing is a legitimate choice, but it is not a free one, and now it has a price on it.

Option 2: Build it yourselves. Every tool named here is off the shelf. Realistically it takes 60 to 90 hours of a senior person's time plus the discipline to measure baselines before and after. The common failure is not technical: the project becomes nobody's job by June. If you go this way, take our tally as your baseline. It is yours to keep either way.

Option 3: Engage nAlte. This assessment costs \$1,500, credited in full against Phase 1 if you engage within 60 days. Phase 1 is Build 1 plus the data policy: \$2,500 build and \$95/mo, so \$1,000 net after the credit. Phases 2 and 3 proceed only if Phase 1 survives its own kill criteria. Three-month initial term, then month to month, cancel any month. Every expected return is written against a tally row before we start, and the kill criteria above go into the agreement as written.

Full program, all three builds: \$13,000 in builds plus \$415/mo in software, about \$17,980 in year one, against \$39,700 addressable. On new cash alone this does not pay back in year one. It pays back when the 1,003 recovered hours are pushed into the backlog and the lapsed renewals, and we will measure whether that redeployment happened.

Prepared by Nathan Flavel, nAlte, Kelowna, BC. Every claim above carries a number or names its owner.