

AI READINESS AUDIT · FULL REPORT

AI Readiness *Report.*

Prepared for Hawthorn Plumbing and Maintenance
12 June 2026 · Confidential

38

COMPOSITE SCORE / 100

Experimenting

Before the score.

The thing that stayed with me from your interview was your description of writing the same five or six quote types fresh every time. That is the business in one sentence. The work is repetitive enough that AI can move it. You already know what AI is good for. And nobody is acting on it yet.

That is not a criticism. It is the honest position of most established businesses your size. Sarah does four to five hours of retyping every week. Dec is showing you AI tools without being asked. ServiceM8 died because nobody owned the rollout. These are signals, not failures.

What follows is structured to be useful. The executive summary is the short version. The action plan is the practical version. The dimension pages are the depth. None of it requires hiring me.

Daniel

DANIEL XUEREBO. DPEXCONSULTING. SYDNEY.

What you're holding.

This report is the output of a fifteen to twenty question adaptive interview, scored against a fixed rubric. Five dimensions, weighted, into one composite score of 38 out of 100. The bands describe what the score means in plain English. Most established small and medium businesses sit somewhere between twenty five and fifty five. The calibration is against reality, not against AI-native software companies.

The report is structured to be useful even if you never speak to DPEX again. If you only have ten minutes, read in this order. One, the note from Daniel. Two, the executive summary. Three, the score page and the headline insight. Four, the first action. Everything else is depth, there when you want it.

Scores are calibrated honestly. The model is explicitly instructed against grade inflation. A thirty here is more useful than a sixty dressed up. Where evidence in your answers was thin, that is noted in the rationale rather than papered over.

CONTENTS

A note from Daniel

01 How to read this report

02 Score and headline insight

03 Executive summary

04 Two scenarios at twelve months

05 Strategy

06 Data

07 Tooling

08 People

09 Governance

10 Prioritised action plan

11 The next quarter, plotted

12 Implementation

13 Scoring methodology

14 Appendix. Interview transcript

WEIGHTS: STRATEGY 25% · DATA 20% · TOOLING 20% · PEOPLE 20% · GOVERNANCE 15%

Where you sit, in one picture.

38

COMPOSITE SCORE / 100

Experimenting

Individuals are trying tools. Nothing is systematic. Value exists in pockets and leaks everywhere else.

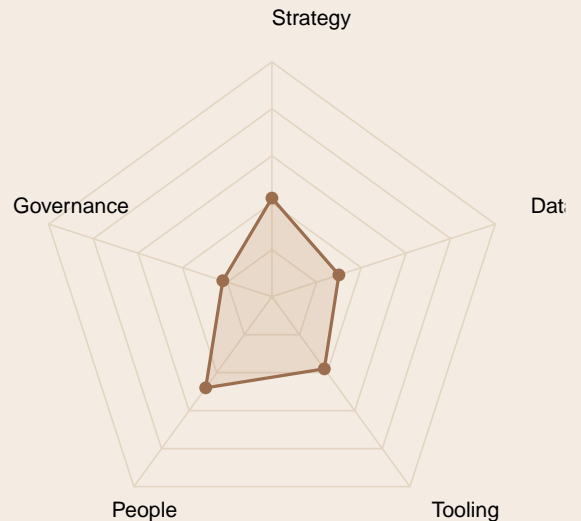


THE HEADLINE

“Your champion exists, your data does not. Fix the foundations before the tooling, in that order.”

RISKS WORTH NAMING

- You would not know if a staff member pasted a client list, body-corporate contact details, or invoice information into ChatGPT this week.
- ServiceM8 was a meaningful spend that died because adoption was treated as the team's problem to solve, not yours.
- Sarah's four to five hours a week of retyping is a single-person-dependency risk as well as a productivity tax.



FIVE-DIMENSION SCORE PROFILE

The reading.

Hawthorn Plumbing and Maintenance sits squarely in the Experimenting band, with a composite of 38. The single clearest thing this audit found is that you already know what AI should fix. Quote turnaround is the explicit pain. You can name the cost in lost work to faster competitors. You have a willing champion in Dec who has shown you tools without being asked. That is a strong starting position.

The gap is foundations. Customer and job data live in three places. Xero, a Google Sheet, and your phone. A four to five hour weekly tax on Sarah for retyping is the visible cost. There is no written rule about what staff can put into AI tools. The team has no shared experience of adopting new software successfully.

The honest reading of an Experimenting score is that you are not behind, and the next move is not buying more software. The next move is cheap, deliberate, and owner led. The action plan that follows orders five moves so the foundations land first.

COMPOSITE

38

Experimenting

Strategy	42
Data	30
Tooling	38
People	48
Governance	22

Act, or do not.

IF YOU ACT ON THIS REPORT

Twelve months from now, quote turn-around is under four hours for the six common job types and you can show the number to anyone who asks. Sarah's retyping hours have halved. Dec runs a thirty-minute Friday review with you where he shows what is working. Your three core systems talk to each other well enough that nobody is opening more than two tabs to write an invoice. The one-page AI policy is up in the office.

IF NOTHING CHANGES

Twelve months from now, quote turn-around still floats between two and seven days because nobody is timing it. ServiceM8's replacement, whatever it is, also dies because the rollout pattern has not changed. Sarah still retypes four hours a week and is one ankle injury away from the back office stalling. Dec has stopped showing you AI tools because nothing happened the last three times he did. A body-corporate client asks you in writing how you handle their data in AI tools. You stall and lose the contract.

WHAT FOLLOWS

The next five pages take each dimension in turn. Score, evidence, strengths, gaps, what would need to be true to move up one band, and the common ways progress on this dimension goes wrong. The action plan after that sequences five moves into the next twelve weeks.

42

Strategy

*A clear outcome, no baseline, no measurement. The ambition is real.
The machinery is not.*


DASHED MARKER: NEXT BAND THRESHOLD (50).

WHAT IS WORKING

You have named the right outcome. Quote turnaround is concrete, customer-visible, and tied directly to lost revenue. You have been realistic about position and budget.

WHERE THE GAP IS

There is no baseline. You estimate two to three days but nobody is timing it. Without a number to beat, any later claim of 'AI is working for us' is unfalsifiable.

TO MOVE UP ONE BAND

Two weeks of simple manual logging of quote send times and acceptance outcomes would produce a baseline. Sarah running that logging, not you, would de-risk the second issue.

WATCH-OUTS

- Treating the quote-time number as a side project of the back office.
- Confusing 'we installed an AI tool' with 'we moved the metric'.

WHAT YOU TOLD ME

“Getting quotes back to customers faster.”

“I would measure it by quote turnaround time and probably win rate.”

30

Data

Records split across Xero, a Google Sheet, and your phone. Sarah is the only bridge.



DASHED MARKER: NEXT BAND THRESHOLD (50).

WHAT IS WORKING

You can name your systems, which is more than most businesses at this stage can do. You have also been honest about the cost.

WHERE THE GAP IS

The customer record has no single source of truth. The Google Sheet is fragile. The phone-as-database means knowledge walks out the door if you do.

TO MOVE UP ONE BAND

One system that holds the customer record, with the others syncing from it. A simple, well-set-up Xero contact module might be enough at your size.

WATCH-OUTS

- Trying to fix all three places at once.
- Doing the migration without naming who owns customer record hygiene afterwards.

WHAT YOU TOLD ME


*“Xero for invoicing and accounts.
The Google Sheet for active jobs.”*

*“Probably four to five hours a week
of just retyping.”*

38

Tooling

Personal ChatGPT use, no embedded workflow, and a software graveyard you cannot fully account for.



DASHED MARKER: NEXT BAND THRESHOLD (50).

WHAT IS WORKING

You are using AI yourself, on a real task, and you have identified the highest-leverage repetition in the business unprompted.

WHERE THE GAP IS

None of it is embedded. Twice-weekly personal ChatGPT use is not a workflow. Nobody else benefits from it.

TO MOVE UP ONE BAND

A single named workflow with AI in it that is used by someone other than you, weekly. The quote-drafting task is the obvious candidate.

WATCH-OUTS

- Choosing a tool before choosing the workflow.
- Skipping the software audit because it is boring.

WHAT YOU TOLD ME

“I use ChatGPT for writing customer emails when something tricky comes up.”

“Writing quotes. Same five or six job types over and over.”

48

People

Dec is your champion. He does not yet know you would back him. That is the highest-leverage gap.

 DASHED MARKER: NEXT BAND THRESHOLD (50).

WHAT IS WORKING

Dec is the asset most businesses at this band do not have. You also know exactly why ServiceM8 died.

WHERE THE GAP IS

Dec does not know he has your backing. The older crew thinks it is a fad. Sarah's concern about losing the personal touch is worth taking seriously.

TO MOVE UP ONE BAND

An explicit, paid hour a week for Dec, owned by him, reporting back to you monthly. Plus one piece of training on a single tool.

WATCH-OUTS

- Promoting Dec to 'tech guy' without protecting his actual tools-on-the-wall time.
- Letting Sarah's 'personal touch' concern get dismissed.

WHAT YOU TOLD ME

“One of my younger guys, Dec, is actually pretty into it.”

“He does not know I would properly back him though.”

22

Governance

Nothing written. Nothing said. No incident path. The cheapest dangerous gap to close.



DASHED MARKER: NEXT BAND THRESHOLD (25).

WHAT IS WORKING

You have not dressed any of this up. You know the policy does not exist.

WHERE THE GAP IS

Three concrete holes. No rule for what staff can put into public AI tools. No way to detect if confidential data is exposed. Body-corporate clients have raised data handling without you having a straight answer.

TO MOVE UP ONE BAND

A one-page written policy covering 'public AI tools and our customer data' that the team sees and signs. Plus a single-sentence answer you can give a body-corporate client.

WATCH-OUTS

- Writing a five-page policy that nobody reads.
- Treating this as a compliance chore rather than a sales-enabler.

WHAT YOU TOLD ME

“Nothing written. Nothing said either really.”

“I would never know unless something went badly wrong.”

The next quarter, in order.

Five moves, sequenced so each unlocks or de-risks the next. None require hiring a consultant or buying enterprise software. The first is startable this week at zero cost.

01. Write the one-page AI use policy

GOVERNANCE . WEEKS 1 . ONE HOUR, OWNED BY YOU, DONE THIS WEEK

THIS WEEK

Open a Google Doc on Monday. Two sections. What staff can and cannot put into public AI tools. What to do if you think the rule was broken. Print it, sign it, get the team to sign it Friday.

WHY NOW

You said you would never know if confidential data got pasted into ChatGPT. That is the cheapest, most dangerous gap to close. The cost of fixing this is forty minutes.

WHAT CHANGES WHEN IT LANDS

You have a one-page document you can hand to a body-corporate client when they ask. Your worst-case AI incident is now detectable.

COMMON FAILURE MODE

Do not let this become a five-page policy.

02. Tell Dec he has your backing

PEOPLE . WEEKS 1 . ONE COFFEE, OWNED BY YOU, DONE THIS WEEK

THIS WEEK

Take Dec for a coffee. Tell him plainly. I want you to be the AI guy. I will pay you one hour a week to try things. You report back monthly.

WHY NOW

He already does this on his own time without being asked. He does not currently know you would back him.

WHAT CHANGES WHEN IT LANDS

Dec runs a monthly thirty-minute review with you. The older crew sees the rollout has a peer doing it.

COMMON FAILURE MODE

Do not promote him out of his trade work to do this. Protect the hour.

03. Baseline quote turnaround for two weeks

STRATEGY . WEEKS 1 TO 3 .

FIFTEEN MINUTES SETUP, TWO WEEKS ELAPSED, SARAH AT TEN MINUTES A WEEK

THIS WEEK

Add two columns to the Google Sheet. Enquiry received and quote sent. Tell the team to fill them in for every quote for two weeks.

WHY NOW

You named quote turnaround as the outcome you most want AI to move, but nobody is timing it. Two weeks of logging gives every later decision a number to beat.

WHAT CHANGES WHEN IT LANDS

You have a real average and median. Sarah owns the number rather than you.

COMMON FAILURE MODE

Do not react to the number for the first two weeks. The point is to measure, not motivate.

04. Audit and cut the software graveyard

TOOLING . WEEKS 2 TO 3 .

HALF A DAY FOR SARAH, SIGN-OFF FROM YOU, THIS FORTNIGHT

THIS WEEK

Sarah pulls the last three months of credit card statements and circles every recurring software charge. By Friday you have a list with three buckets.

WHY NOW

You said you would have to check the statement to know what software you are paying for. The cancellations free up the budget you already nominated for AI tooling.

WHAT CHANGES WHEN IT LANDS

You know exactly what software you are paying for. The next AI tool comes out of zombie-subscription savings.

COMMON FAILURE MODE

Do not cancel the 'unsure' bucket. Pause those for two months.

05. Build one AI-assisted quote template

TOOLING . WEEKS 5 TO 8 . HALF A DAY ACROSS DEC AND YOU, THIS MONTH

THIS WEEK

Pick the most common of your five or six job types. Dec and you draft an AI prompt that turns a short job description into a quote in your house voice. Test it on three real recent jobs.

WHY NOW

This is where the quote-turnaround number actually moves. It needs to come fifth because it depends on the other four landing first.

WHAT CHANGES WHEN IT LANDS

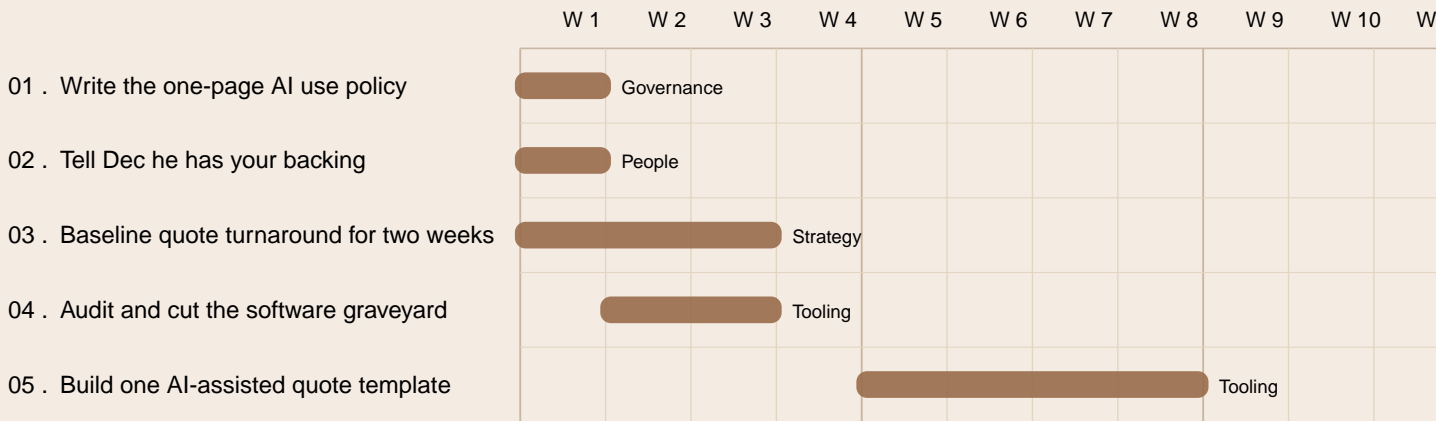
One job type drafts faster. The team sees AI doing real work, not theoretical work. The baseline number moves.

COMMON FAILURE MODE

Do not bypass Sarah on the voice. Her concern about losing the personal touch is the real risk here.

Twelve weeks. Five moves.

Each bar shows when an action runs. Sequencing matters. The foundations land before the AI moves rest on them.



WEEKS ARE CALENDAR WEEKS FROM THE DAY THIS REPORT ARRIVES. TREAT THEM AS PLANNING GUIDES, NOT CONTRACTS.

When you want this done, not just diagnosed.

This report tells you where you stand and what to do first. If you want a hand actually doing it, DPEX runs fixed-scope implementation engagements that take one action from this list and embed it in your business alongside your team.

The shape. Two to six weeks, scoped against the specific action and outcome you want, delivered in your tools and on your cadence. The team owns the result after I leave. No retainer. No account-management layer. You work with me.

To book a thirty-minute discussion: reply to the email this report arrived with, or write to dpexconsulting@gmail.com with the subject "Implementation". I read every message myself and reply within two business days. No funnel. No sales pressure. If we're not the right fit I'll say so.

Daniel Xuereb

DPEXCONSULTING · SYDNEY

ABOUT THIS REPORT

Generated by DPEXConsulting's Readiness Audit engine, scored against a fixed rubric and reviewed against behavioural anchors per question. The on-screen result and this PDF carry identical scores; this PDF is yours to keep, share, and re-read.

DISCLAIMER

This audit is a structured diagnostic, not legal, financial, or cybersecurity advice. The actions are starting points; sector regulation may layer additional requirements. The score is a calibration tool, not a guarantee of outcome.

How the score was built.

The five dimensions

Each dimension carries a 0 to 100 score, behind which sits a set of anchored interview questions. The questions probe evidence, not intention. The composite is the weighted average shown opposite. Strategy carries the most weight because it sets the ceiling on what the others can achieve.

How a score is assigned

Each answer is scored against a calibrated anchor. High scores require named tools, named people, numbers, or specific examples. Generic answers and vague intentions land in the lower half of the range. The grading explicitly penalises optimism. The model is instructed that an honest thirty is more useful than a kind sixty.

The maturity bands

0 to 24	Observing	AI is something happening to other businesses. The opportunity is real but nothing is in motion yet.
25 to 49	Experimenting	Individuals are trying tools; nothing is systematic. Value exists in pockets and leaks everywhere else.
50 to 69	Operationalising	AI is in at least one real workflow with someone accountable for it. The foundations can now compound.
70 to 84	Scaling	AI is embedded across several functions with measurable returns. The work now is depth and governance.
85 to 100	AI-native	AI is part of how the business runs and decides. Few businesses score here honestly.

DIMENSION WEIGHTS

Strategy	25%
Data	20%
Tooling	20%
People	20%
Governance	15%

WHAT THE SCORE IS NOT

The score is a structured diagnostic, not legal, financial, or cybersecurity advice. Sector regulation may layer additional requirements on the actions proposed. The score is a calibration tool, not a guarantee of outcome.

What you told me.

Every question asked, in order, with your answer. The score above was derived from these answers. Kept here verbatim so you can revisit any line without having to remember the conversation.

Q1 · STRATEGY

What is the one business outcome you would most want AI to move this year?

Getting quotes back to customers faster. We lose work because we take two to three days to quote and our customers go cold or call the next guy. I would measure it by quote turnaround time and probably win rate, but honestly we do not track win rate at the moment.

Q2 · STRATEGY · FOLLOW-UP

What is the current average quote turnaround, in hours?

Honestly couldn't tell you exactly. Feels like two to three days but some sit a week. Nobody is timing it.

Q3 · STRATEGY

Where does AI sit on your priority list right now?

Probably third or fourth. Top is hiring a second tradesman, then getting the back office under control, then this. Budget is realistic, maybe three to five hundred a month for tools.

Q4 · STRATEGY

Who owns the question of what to do about AI?

Me. Along with everything else. My wife Sarah does the books and would have ideas too but it has been on my plate.

Q5 · DATA

If you needed your last twelve months of customer or job data in a spreadsheet today, how long would that take?

A couple of days. Sarah would do it. Some of it is in Xero, some in our job sheet which is a Google Sheet, and a bunch of customer details are in my phone or emails. We do not have one system.

Q6 · DATA

Walk me through where your core business records actually live.

Xero for invoicing and accounts. The Google Sheet for active jobs and scheduling. Customers are everywhere, phone contacts, Gmail, written down on the invoices. We tried ServiceM8 a couple years ago and did not stick with it.

Q7 · DATA

What information does your team retype from one place into another every week?

Job details from the sheet into Xero when we invoice. Customer details when a repeat customer calls and we look them up. Quote details into emails. Probably four to five hours a week between Sarah and me of just retyping.

Q8 · TOOLING

What AI tools does anyone in the business use weekly right now?

I use ChatGPT for writing customer emails when something tricky comes up, maybe twice a week. That is it. None of the lads use anything.

Q9 · TOOLING

What are you currently paying for in software that nobody really uses?

ServiceM8 we cancelled. Definitely paying for Microsoft Office licences nobody opens. Probably some Adobe thing. I would have to check the credit card statement.

Q10 · TOOLING

What is the most repetitive task in the business?

Writing quotes. Same five or six job types over and over. I have rough templates in my head but I type them fresh every time.

Q11 · PEOPLE

When AI comes up with your team, what is the honest reaction?

Mixed. One of my younger guys, Dec, is actually pretty into it. He has shown me a few things he has tried on his phone. The older guys think it is a fad. Sarah is curious but worried about losing the personal touch with customers.

Q12 · PEOPLE

Who would be first to make an AI tool actually work?

Dec, for sure. He does not know I would properly back him though. We have not had that conversation. I should probably tell him.

Q13 · PEOPLE

When you have rolled out something new before, how did the team learn it?

They figured it out. Which is why ServiceM8 died. Nobody invested in showing them how it worked, so they kept using the sheet.

Q14 · GOVERNANCE

Is there anything written down about what staff can put into AI tools?

Nothing written. Nothing said either really. I do not think the team would put customer details into ChatGPT but I have not told them not to.

Q15 · GOVERNANCE

If a client confidential detail ended up in a public AI tool tomorrow, how would you find out?

I would never know unless something went badly wrong. Wouldn't have a clue what to do, call the lawyer I suppose.

Q16 · GOVERNANCE

Would AI use change your standing with body-corporate clients?

Hadn't really thought about AI changing that. We do some work for body corporates and they sometimes ask about data handling.