SAMPLE DELIVERABLE · MGT-09

AI Augmentation Business Plan

Augment, not replace: where assistance genuinely moves hours, funded in stages against measured effect, with every consequential output on a human-owned pathway.

Prepared for Merrow & Co (fictional) by Red Cell Consulting

CONTENTS

About this document

This is a sample of the deliverable Red Cell Consulting produces in a real engagement. Every organisation, figure and finding in it is fictional, and the structure, depth and standards are exactly what a client receives. Where a real engagement would cite source documents, the sample says so.

Executive summary

Merrow & Co is a ninety-person professional services firm. This plan moves roughly 750 hours a month from assembly work to judgement work within a year, for a net year-one effect of about A$258k, gated so the spend follows the evidence.

The thesis

The gains in a firm like Merrow are not in replacing professionals; they are in removing the assembly work wrapped around professional judgement: first passes, summaries, drafts, notes and formatting. AI does assembly well under supervision. Judgement, advice and sign-off stay human, by design rather than by hope.

Why a business plan and not a pilot

Unstructured pilots produce enthusiasm and no evidence. This plan baselines the hours honestly, stages the rollout behind gates, prices the effort truthfully (including internal time), and gives the board a kill-switch it will never need to debate: gates either pass or they do not.

PLAN ON A PAGE

ELEMENT	POSITION
Workflows in scope, year one	Five, selected from a measured two-week baseline of where hours actually go
Hours effect at target	~750 hours/month from month six, measured monthly, banked only when measured
Investment	A$154k year one (licences A$96k, change effort A$58k), stage-gated
Net year-one effect	A$258k; break-even month eight if gates hold
Non-negotiables	Client data stays in tenancy; consequential figures from systems of record; human sign-off on all advice

Method: how the plan was built

Two weeks of measurement before a single tool was named. The plan is only as good as the baseline underneath it.

THE BASELINE

- Time captured across the firm for a fortnight against a twelve-activity frame, calibrated with partners first so the categories match how work actually happens.
- Interviews with eleven staff across levels: where the friction is, what gets done at 9pm, what juniors actually spend their first year doing.
- Document flows sampled: how a typical engagement letter, report and proposal actually get assembled, versioned and reviewed.

WHAT THE BASELINE SHOWED

FINDING	IMPLICATION
38 percent of professional hours are assembly: reading-to-summarise, drafting-from-precedent, formatting, note-taking	The augmentation target is large and specific, not speculative
Senior staff do 21 hours/month of work juniors could do with better tooling	Leverage problem, not headcount problem
Three staff already use consumer AI tools quietly, outside any policy	Shadow usage exists; governance is overdue regardless of this plan
The document management system is the real system of record and its search is poor	Retrieval quality gates everything; fix it first

Selection criteria for workflows

A workflow entered scope only if: hours were measured and material; outputs are reviewed by a human as part of normal practice; the data involved can stay inside the firm's tenancy; and failure modes are visible before they reach a client. Five workflows passed. Several popular ideas did not, and page four says why.

Task analysis: where the hours are

Five workflows, measured hours, targeted assistance, and the consequence class that decides the controls.

THE AUGMENTATION MAP

WORKFLOW	HRS/MO TODAY	ASSISTANCE	TARGET	CONSEQUENCE CLASS
Document review and summarisation	640	Assisted first-pass with linked citations to source pages	380	Reviewed output
Client reporting assembly	420	Drafting from structured data; human sign-off unchanged	250	Reviewed output
Research and precedent search	310	Retrieval with sources; unsourced answers disabled	200	Advisory input
Meeting notes and follow-ups	260	Transcription, action extraction, draft follow-ups	120	Low
Proposal first drafts	180	Template-grounded drafting from the pricing sheet	110	Reviewed output
Total	**1,810**		**1,060**	~750 hrs/mo effect

WHAT WAS CONSIDERED AND EXCLUDED

IDEA	WHY NOT, FOR NOW
Client-facing chatbot	Consequence class too high for year one; no review step between output and client
Automated advice drafting end-to-end	Advice is the judgement work; assistance drafts inputs to it, not it
HR screening assistance	High consequence to individuals; deferred until the governance layer is proven

Hours from the measured baseline; targets assume month-six adoption. Nothing is banked before it is measured, and the ROI model on page seven uses measured effect only.

What stays human, and the tool policy

The controls are the product. Assistance without boundaries is shadow IT with better marketing.

HUMAN BY DESIGN

- All client advice, pricing and sign-off. Assistance drafts; people decide and sign.
- Anything touching a regulated duty routes through the named owner of that duty.
- Consequential figures come from systems of record, cited; a model's memory is not a source.
- Client communication goes out under a person's name after that person has read it. No exceptions, including the busy weeks, which is when the rule earns its keep.

TOOL POLICY · THE SHORT VERSION THAT FITS ON A CARD

RULE	PRACTICE
Approved tools only	The register lists them, their data flows and their owner; anything else is a request, not a workaround
Client data stays in tenancy	No client material into consumer tools; supplier terms reviewed against this rule before approval
Cite or it does not ship	Assisted output carries its sources; uncited claims are removed at review
Mark the assistance	Internally, assisted drafts are flagged so reviewers know where to look hardest
Report the weird	Confident nonsense, data appearing where it should not: logged the day it is seen

Shadow usage, handled honestly

Three staff already use consumer tools. The response is not discipline; it is a better sanctioned alternative plus a clear rule. Shadow usage is demand signal, and this plan is the supply.

Rollout: stages and gates

Four stages over twelve months. Spend follows evidence through gates agreed now.

STAGE	WINDOW	SCOPE	GATE TO PROCEED
1	M1-2	Document review + meeting notes, one team each; register and policy live from day one	Measured hours effect in range; error rate inside tolerance; zero data-boundary incidents
2	M3-5	Reporting assembly + research; retrieval fix on the document system	Adoption above 60 percent in pilot teams; zero unreviewed outputs shipped
3	M6-9	Firm-wide on proven workflows; proposals pilot	Capacity effect visible in utilisation and turnaround data
4	M10-12	Optimisation; supplier consolidation; year-two plan	Board review against this plan's ROI model

GATE MECHANICS

- Gates are measured by finance and operations jointly, not by the project's own enthusiasm.
- A failed gate pauses spend, names the failing element, and offers the board fix-or-stop with numbers. There is no third meeting.
- The measures are defined now (page six of the engagement version carries the formulas), so nobody negotiates definitions in month five.

Why stage one is boring

Notes and document review are deliberately unglamorous: high hours, low consequence, fast feedback. The goal of stage one is not transformation; it is evidence and habits, the two things every later stage spends.

Capability, governance and the ROI model

Training built per workflow, governance wired from day one, and a model that only counts what is measured.

CAPABILITY AND ADOPTION

- Two-hour hands-on training per workflow, built on Merrow's real documents, not vendor demos.
- Each team names a lead user; their patterns become the team playbook, refreshed monthly.
- Managers review assisted work weekly for the first month of each rollout, then by sample.
- Adoption is measured by usage within the workflow, not logins: a login is not a behaviour.

GOVERNANCE TIE-IN

- Every tool enters the AI system register before first use: purpose, data flows, supplier, named owner, consequence class.
- Evidence kept from day one: baselines, gate results, incidents and decisions, in one file.
- Aligned to the incoming Australian Standards direction and ISO 42001, so nothing here is redone when the rules land in 2027.

ROI MODEL · MEASURED EFFECT ONLY

LINE	YEAR 1	NOTE
Licences and platform	A$96k	Stages 3-4 committed only if gates pass
Training and rollout effort	A$58k	Internal time priced at cost, honestly
Hours effect at target	A$412k	750 hrs/mo at blended rate, from month six
Net year-one effect	**A$258k**	Break-even month eight

ROI SENSITIVITY	NET EFFECT
Adoption reaches 70 percent of target hours	A$135k
Adoption at target, licences 20 percent higher	A$239k
Stage 3 deferred a quarter by gate	A$156k

The capacity question, what the freed senior hours are spent on, is a management decision the plan surfaces monthly rather than assumes.

About Red Cell Consulting

Consulting and advisory, backed by platforms we design, build and operate ourselves.

Red Cell Consulting is a multi-disciplinary consultancy working across enterprise and technology, financial analysis, clinical psychology and behavioural science, advisory and investment readiness, restructuring, management consulting and AI governance. The distinctive part is that we also build: our platforms, including Recursys for deal intelligence, Myra for deterministic financial analysis and ADMRL for structured clinical reporting, run in production under the same disciplines we recommend to clients. When we advise on evidence, testing and governance, we are describing how our own systems work.

HOW AN ENGAGEMENT RUNS

STEP	WHAT HAPPENS
1 · Conversation	A free first discussion of the situation. We will tell you plainly whether we can help and what it would involve.
2 · Scope	A short written scope: the question, the deliverables, the timeline, the fee. Fixed where the work allows it.
3 · The work	Delivered in the structure this sample shows, with working sessions rather than big reveals. You see findings as they form.
4 · Handover	Deliverables that stand on their own, and the option of standing support where it makes sense.

FOR THIS SERVICE

Management consulting at Red Cell now includes exactly this: AI augmentation planned like an operating change, not a gadget rollout. Measured baselines, staged spend, human-owned pathways and governance that will satisfy 2027. If your firm is accumulating tools without a plan, start with the conversation.

Contact

+61 3 9448 4933 · Mon-Fri 9am-5pm Melbourne

info@redcc.au

www.redcc.au