CSDR

PLATFORM

CENTRE FOR SUSTAINABLE DEVELOPMENT REFORM

AI-Native Transition Guide

CLASSIFICATION
INTERNAL

PROJECT
platform

DATE
7 June 2026

Revision 9b000f6 · Built 7 June 2026

AI-Native Transition Guide

INTERNAL

1. Why This Is Happening

AI agent capability is increasing exponentially. Tasks that required significant human effort a year ago are now routine for an AI worker. This trajectory is accelerating, not levelling off.

Well-organised, machine-readable data compounds this capability. Every note, transcript, meeting record, and research file makes every AI worker more capable — multiplicatively, not linearly. A disorganised file share does not compound; a structured information base does. The difference between ad hoc AI use and a consolidated platform is the difference between a tool and infrastructure.

This unlocks a fundamental shift in how work happens. Staff move from doing to orchestrating — directing teams of highly capable AI workers with deep project knowledge, each able to spawn copies of themselves for parallel tasks. A single researcher directing multiple AI workers across different workstreams is not a future aspiration; it is what the platform enables now.

The human work that matters most — stakeholder engagement, creative thinking, relationship-building, strategic judgment — gets more of your time. The production work that currently swamps it gets automated. This is not a technology upgrade. It is a different way of working.

2. What This Is

A proposal for the Centre to move to a new platform in April 2026. This document explains what changes, what stays the same, and how staff would be supported through the transition.

> Ready to try it? The Quickstart Guide walks you through your first tasks step by step — record a meeting note, request a brief, and download branded outputs. No prior experience required.

The platform consolidates the Centre's information base — research, notes, transcripts, project files, meeting records — in one place. Publications (reports, briefs, presentations, client deliverables) are generated from this base and deployed automatically to different audiences. Binary files (images, charts, videos, PDFs) are stored separately in cloud storage and linked seamlessly. Email, calendar, and Slack remain unchanged, but the Centre's substantive work — everything that contributes to projects — lives in the platform.

Each staff member is complemented by a cohort of dedicated AI worker(s). The substantive capability of large language models has now surpassed the vast majority of technical specialists for most analytical tasks across law, economics, environmental science, and policy. This is not a hedge or a projection — it is a measured reality as of early 2026. Your role moves up a level of abstraction: you guide, enhance, and apply judgment rather than produce from scratch. This is not a diminishment; it is what makes a team of 20+ staff with 20+ dedicated AI workers — each able to spawn copies of themselves as needed for parallel tasks — fundamentally more capable than any traditional research centre.

The platform deploys content to three audience tiers:

- The internal site (internal.csdr.global) — the full information base, visible to all staff.
- The public research site (docs.csdr.global) — content marked for external audiences.
- Client-specific sites (e.g. alpha-review.csdr.global) — individual project sites with email-gated access for specific clients. Each partner gets a dedicated website with live, always-current reporting. No emailing attachments, no version confusion. Partners see the current state of their project at any time, gated by email access. Updates publish automatically when approved.

All three are generated from the same source content. A single metadata field (`visibility`) controls which audiences see what. You do not need to produce separate versions for different audiences — the platform handles that.

3. How Work Changes

The way staff produce work changes fundamentally. Less time drafting, formatting, and managing documents. More time guiding, reviewing, and engaging. This section is honest about the scale of the shift — and the practical sections below show exactly what it looks like.

What genuinely stays the same

- Slack and email — for internal and external communication, largely as they are now. The change: Slack discussions that inform project decisions should be summarised in the platform, where AI workers can access them. Not every Slack message needs to move — only substantive observations, decisions, and context that contribute to a project. Post links in Slack for discussion; keep the substantive record in the platform. Future: AI workers may be able to ingest Slack channels directly, but this is a roadmap item, not a current capability.
- Calendar and meetings — unchanged.
- Stakeholder relationships and engagement — unchanged. This is the work the Centre exists to do.
- Your expertise and judgment — you decide what is worth saying, whether the analysis is sound, and whether the output meets the Centre's standards.

What changes

- Production shifts from writing in a word processor to directing AI workers. This is not a case of typing one sentence and receiving a finished document. Good specification requires thought — what is the purpose, who is the audience, what structure serves the argument, what context does the AI worker need. The quality of what you get out is determined by the quality of what you put in. This is design thinking: you are the architect, your AI worker is the builder.
- Review shifts from tracked changes in emailed Word files to rendered output on the internal site. You see the finished product — web page, PDF, DOCX — and request changes conversationally.
- Iterative refinement is the norm. A first draft is a starting point. You review, redirect, add context, sharpen the argument across multiple rounds. This is not a failure of the AI — it is how the workflow is designed. Each round gets better as the AI worker accumulates your feedback and intent.
- Notes and project records move into the platform where AI workers draw on them. Meeting notes, observations, research fragments — everything that contributes to a project belongs in the information base. The more context your AI worker has, the better it performs.
- Personal scratchpads remain your business, in whatever tools you prefer. But project-related material belongs in the platform.

The shift is real, but supported at every step. The practical sections below show exactly what it looks like day to day.

4. What Changes — Before and After

The shift is from manually producing documents to specifying what you need and reviewing the result — from doer to orchestrator.

	Traditional approach	Platform approach
Producing content	Write and format in Word from scratch	Describe what you need; your AI worker produces a first draft
Research and analysis	Compile manually from multiple sources	AI worker draws on the full project information base — every note, transcript, and file
Formatting	Manual template wrestling	Automatic — branding, layout, styles handled by the build pipeline
Review	Email files, tracked changes	Review rendered output on the internal site; request changes conversationally
Multi-format output	Save as PDF, reformat for each audience	PDF, DOCX, and web pages generated automatically from one source
Distribution	Email attachments, manage versions	Publish once; platform deploys to all relevant sites and partner sites
Version control	”Final_v3_ACTUAL_FINAL.docx”	Full history tracked automatically; any version restorable
Project knowledge	Scattered across emails, drives, personal notes	Consolidated in the platform; AI workers have full context
Quantitative analysis	Build spreadsheets, run Stata manually	Describe the analysis; AI worker produces code, runs it, generates tables and charts

The platform handles production and distribution. You focus on substance.

5. The Tools You Will Use

You need two things on day one: Claude Co-Work `[https://claude.ai]` (a desktop app where you describe what you need in plain English) and the internal site (internal.csdr.global, where you browse and review all Centre content). Everyone gets a GitHub `[https://github.com]` account — set up for you during orientation — which is how you log in and how changes are tracked.

For staff who want more direct control, Claude Code `[https://claude.ai/claude-code]` in VS Code `[https://code.visualstudio.com]` is the most capable way to work with the platform. It is not the starting point — staff progress to it at their own pace, if they choose to. The progression is:

1. Claude Co-Work + internal site — where everyone starts.
2. Claude Code in VS Code — the ideal destination for those who want it, but not a requirement.

The Quickstart Guide walks you through setting up and completing your first tasks step by step.

6. What Your Day Looks Like

Everything the Centre produces is for a specific audience and purpose. Understanding who you are producing for shapes how you work with the platform:

Audience	Typical outputs	Distribution channel
Minister's office / senior government	Policy briefs, executive summaries, cost-benefit analyses	Branded PDF emailed or via client spoke site
Donor / funding body	Progress reports, compliance reports, acquittals	Branded DOCX in their template, or client spoke site
External research partners	Co-authored reports, working papers, datasets	Public research site, or Google Docs for collaboration
Client / project partner	Assessment reports, technical analyses, deliverables	Dedicated client spoke site (email-gated, always current)
Internal (Centre staff)	Meeting notes, project records, research fragments, dashboards	Internal site (visible to all staff)
General public	Published research, media summaries, public briefs	Public research site

Your AI workers produce artefacts for these audiences. The platform handles formatting, branding, and distribution to the right channel. Your job is to specify what the audience needs, review the output, and approve.

Producing a policy brief

1. Open Claude Co-Work. Start by planning: “I need a policy brief on renewable energy procurement for the climate-risk project. The audience is the minister’s office — they need cost-benefit framing, not technical detail. Pull from the project’s energy analysis files and the stakeholder meeting notes from February. Structure as context, findings, three recommendations, and a summary table with cost estimates.”
2. Claude produces a first draft plan, the staff member iterates conversationally, then executes drafting followed by Claude submitting for review.
3. Review on the internal site. The structure works but the recommendations are too generic — they need to reference the specific procurement framework the minister’s office uses. Back in Claude Co-Work: “The recommendations need to reference the NSW Renewable Energy Procurement Framework specifically. Recommendation 2 should focus on battery storage co-investment — use the figures from the Alpha Corp assessment. Also, the tone is too academic for this audience — more direct, fewer caveats.”
4. Claude revises. Review again. The substance is right but the summary table needs a column for implementation timeline.
5. “Add an implementation timeline column to the summary table. Q3 2026 for recommendation 1, Q1 2027 for recommendations 2 and 3. Approve when done.”
6. The platform generates a branded PDF, DOCX, and publishes the web page.

The point: you are designing the output, not just requesting it. Each round sharpens the result because you bring judgment and context that the AI worker cannot generate on its own. The investment in clear specification pays back multiplicatively — a well-specified brief takes three rounds; a vague one takes ten.

Sending a report to a client

Two options:

- Download and email — go to the internal site, download the branded PDF or DOCX, and email it as you normally would.
- Client site — if the client has a dedicated spoke site, they access it directly with their email address. You do not need to send anything; the content is already there.

Recording project information

After a stakeholder meeting, open Claude Co-Work and tell it what was discussed. It creates a structured record in the project. Your AI worker draws on this — along with every other note, transcript, and file in the project — when drafting the next report or brief. The more context in the platform, the better your AI worker performs.

Updating existing content

Open Claude Co-Work: “In the Alpha Corp assessment, update the energy storage section. Battery capacity is now at 45 per cent of optimal across all facilities, up from 30 per cent.” Claude finds the document, makes the change, and submits it for review. You check it on the internal site

and approve.

Running a quantitative analysis

1. Open Claude Co-Work: “I need to analyse the quarterly financial data for the GOAP portfolio. The raw data is in the project — it’s the UNSW finance download from last week. I want variance analysis against budget by project, a summary table, and a chart showing spend trajectory against forecast for each active project. Use Python.”
2. Claude produces analysis code, runs it against the data, and generates the outputs — tables and charts ready to embed in a report.
3. Review the results. “The variance calculation for Project Delta looks wrong — it should be against the revised budget from November, not the original. Also add a year-to-date cumulative column.”
4. Claude corrects and re-runs. The updated analysis is version-controlled — next quarter, run the same analysis against new data with one instruction.

This is a shift from building spreadsheets to describing what you need analytically. Your expertise — knowing what analysis matters, what the numbers mean, where the figures should come from — remains the valuable part. The code is the AI worker’s job. Claude Co-Work or Claude Code can produce analysis in R, Python, or JavaScript depending on what the task requires. The results are reproducible, auditable, version-controlled, and embed directly into reports and publications.

For staff who currently work in Excel or statistical packages like Stata: you do not need to learn to code. You need to learn to describe your analytical requirements clearly — which is what you already do when explaining your analysis to a colleague. The AI worker handles the implementation.

The Quickstart Guide walks you through these tasks step by step. This section gives the shape of a typical day; the quickstart gives the exact steps.

7. Common Tasks — Office 365 to Platform

What you need to do	How you do it now	How you do it on the platform
Make a quick table	Format in Word	Tell Claude or type simple markdown
Email a PDF to a minister	Create in Word, save as PDF, email	Download branded PDF from internal site, email as normal
Find last quarter's report	Search email, shared drives	Fuzzy and keyword search instantly through the website, or deep search using Claude Co-Work or Claude Code
Share a draft with a colleague	Email a Word file	Point them to the page on the internal site
Collaborate with an external partner	Email Word files back and forth	Collaborate in Google Docs; download as markdown for production; Claude converts and submits for review
Give a client access to a deliverable	Email files or manage SharePoint	Add their email to the content metadata; access provisioned automatically
Create a presentation	Build slides in PowerPoint	Describe the presentation to Claude; review slides on the internal site
Something goes wrong with AI output	N/A	Nothing publishes without your approval. Request corrections, review again.
System is down and you need something urgently	N/A	Word and email still work. Use them for emergencies. The platform is not a single point of failure for all communication.

8. What the Platform Handles Automatically

Once content is approved, the platform takes care of:

- Formatting and branding — CSDR print CSS for PDFs (A4, correct margins, headers, footers, page numbers, Source Serif 4 body text). Pandoc reference templates for DOCX. Consistent across every output.
- Multi-format generation — one markdown file produces a web page, a branded PDF, and a branded DOCX simultaneously.
- Multi-audience deployment — the same content serves different audiences via different sites. The internal site shows everything to staff. The public site shows content marked as `visibility: public`. Client spoke sites show only that client's content, gated by email. All controlled by a single metadata field — no separate versions to maintain.
- Version control — every change tracked with full history. Any version restorable. Complete audit trail of who changed what, when, and why.
- Access control — driven by content metadata, provisioned automatically via Terraform and Cloudflare Zero Trust. Adding or removing a client's access is a metadata change, not an IT request.
- Sidebar and navigation — the internal site's sidebar is auto-generated from content metadata. New projects appear in the right place automatically.
- Binary asset management — images, charts, and documents stored in R2 cloud storage with a searchable index maintained in the platform. Claude Co-Work handles uploads seamlessly; more technical users use wrangler directly. Your AI worker can search and reference any asset in the index.

The infrastructure follows the content. Adding a new client project with email-gated access, branded PDF output, and a dedicated site requires creating a content file with the right metadata and approving it. Everything else is automatic.

9. Communications and Dissemination Pipeline

The platform automates most of the communications pipeline. Understanding who controls what — and where outputs go — is important for both research staff and the communications team.

How content reaches different audiences

Audience	Channel	How it works
Centre staff	Internal site	All approved content appears automatically. No action required.
General public	Public research site	Content marked `visibility: public` publishes automatically.
Client / project partner	Client spoke site	Content marked with the client's project appears on their dedicated, email-gated site. Updates publish automatically when approved.
Government / minister's office	Branded PDF or DOCX via email	Download from the internal site and email as normal.
Donor / funding body	Branded DOCX (often in their template) or client spoke site	Claude produces formatted outputs in the required template.
Broader stakeholder networks	Slack, email, social media	Share links to the public site or attach branded PDFs.

Who controls what gets published where

- The `visibility` metadata field determines which sites display the content (internal, public, client). The content author sets this; it can be changed at any time.
- The content author approves submissions. Nothing publishes without explicit approval.
- The communications team designs the templates, branding, and formats that the platform applies automatically. They do not need to produce individual outputs — the platform handles that — but they control the design system that ensures consistency across all publications.
- Project leads decide when content is ready for external audiences and which distribution channels to use.

External dissemination workflow

For content that needs active distribution (not just publication on a site):

1. Approve the content on the platform. Branded outputs (PDF, DOCX, web page) are generated automatically.
2. Download the appropriate format from the internal site — PDF for formal distribution, DOCX if the recipient needs an editable version.
3. Distribute through normal channels — email to stakeholders, share links on Slack, post to social media, or upload to SharePoint for UNSW record keeping.

The platform produces the outputs; you control the distribution. For client spoke sites, distribution is automatic — partners see updates as soon as content is approved.

10. Working with External Collaborators

External collaborators use Word and Google Docs. They will continue to do so. The platform does not require them to change.

Google Docs remains the external collaboration platform. Collaborate with partners in Google Docs as you do now. When content is ready for production, give the content to Claude for conversion and submission. Documents do not go to production from a word processor except in exceptional circumstances — the platform is the production system; Google Docs is the collaboration surface.

For detailed workflows covering co-authoring, receiving tracked changes, submitting in external templates, and managing multiple Word versions, see the Collaboration Guide.

11. What the Platform Makes Possible

When every staff member has an AI worker with full project context and the ability to spawn copies for parallel work, a range of tasks become routine that would previously have been impractical:

- Automated donor reporting — AI workers produce complete reports per donor on schedule, drawing on the project information base. Staff review and approve; the platform publishes and distributes.
- CRM-style stakeholder reporting — automatic generation of internal reference reports per stakeholder, consolidating all interactions, project involvement, and engagement history.
- Financial reporting and analysis — bulk data downloads from UNSW systems as input; AI workers produce analysis, variance reports, and summaries for review.
- Impact reporting — automated alignment of project outputs to the UNSW Societal Impact Framework, drawing on project data across the Centre.
- GEDSI review and planning — AI workers review all relevant GOAP projects for gender equality, disability, and social inclusion considerations, producing consolidated assessments and planning recommendations.
- Autonomous research supplementation — AI workers continuously scan across projects for cross-cutting themes, connections, and opportunities that no individual team member has identified. Unseen opportunities surfaced automatically.
- In-country political economy — for in-country staff, the platform enables strengths to compound from sustained information flows on local political economy. Meeting notes, observations, relationship maps, and contextual analysis accumulate in the platform. AI workers draw on this growing base to produce increasingly informed analysis — closer to a organisational structure of a journalism or intelligence agency than a traditional research centre. Local knowledge compounds rather than dissipating.
- Mapping and data visualisation — code that generates maps and visualisations can be produced once and adapted across deliverables, drawing on the spatial data framework. Instead of manually reproducing a map for each report, the code runs against updated data and the output embeds directly. One mapping product serves multiple publications.
- Reproducible quantitative analysis — financial models, statistical analysis, and data processing that currently live in spreadsheets move to version-controlled code. AI workers produce the analysis code from staff specifications; the same analysis re-runs against updated data each period. Results embed directly into reports. One analytical framework serves multiple reporting cycles.

These are not hypothetical. They are the operational target for the platform — the kinds of tasks that become routine when the information base is consolidated and AI workers have full context.

12. What Is Genuinely Harder

Honesty about what the transition costs:

- The first two weeks are unfamiliar. Conversational interfaces feel different from menus and toolbars. This is a real adjustment, not a trivial one.
- Complex bespoke formatting takes more effort than dragging boxes in PowerPoint. The platform handles standard layouts well; unusual one-off designs require more work.
- External collaboration adds a conversion step. Working in Google Docs and bringing content into the platform for production is functional but not invisible. Google Docs remains the working surface for live co-editing; the platform takes over for production.
- You need a GitHub account and need to learn to log in via GitHub. This is set up for you, but it is new.
- You are learning a new way of working. The effort is real. It is a one-time cost, but it is a cost, and pretending otherwise would be dishonest.
- Internet access is required. Claude Co-Work, Claude Code, and the internal site all require an internet connection. For in-country staff or low-connectivity environments, this is a real constraint. Word and email remain available offline as a fallback.

These are known trade-offs, not surprises. Support is provided specifically because the adjustment is real.

13. How You Will Be Supported

Week 1: Orientation

- GitHub accounts set up for everyone. Claude Co-Work installed on your machine.
- Browse the internal site in read-only mode. See how projects are structured. See your colleagues' work rendered as web pages, PDFs, and Word documents.
- No expectations to produce anything. The goal is familiarity.

Week 2: Paired production

- You describe what you need — a report, a brief, a presentation — verbally or in writing.
- A colleague or Claude produces it. You review the output on the internal site, request changes, and approve.
- By end of week two, you have seen your own work published through the platform.

Ongoing support

- A dedicated Slack channel for platform questions, solutions, and tips.
- Individual support for anyone who needs it, at their own pace.
- The platform and wider AI-native operations would benefit significantly from a full-time junior software developer on the operations team to provide dedicated platform support and development. This person would handle troubleshooting, build issues, first-line support for staff learning the workflow, and process improvement. This would be a structural addition to the team, not a temporary measure, replacing highly manual processes in finance, contracting, admin and data management (which consumes significant amounts of substantive staff time) to the maximum extent possible.
- A 2-hour weekly drop-in clinic for platform questions, hands-on help, and troubleshooting. Open to all staff, no appointment needed. The clinic runs weekly until it is no longer necessary — the pace of the transition determines the schedule, not an arbitrary end date.

“Specify and review” is a valid permanent workflow. You describe what you need; Claude or a colleague produces it; you review and approve. You never need to learn git, markdown, or the command line. This is not training wheels — it is a complete, professional way of working that many staff will use indefinitely.

14. Skill Levels — Go at Your Own Pace

Level	Tools	What you do	Who this suits
Specify and review	Claude Co-Work + internal site	Describe what you need, review output, approve	Everyone, from day one
Direct editing	Claude Co-Work + browser editing	Make simple changes to content directly	Curious staff wanting more control
Independent agent use	Claude Code in VS Code	Use AI agents to produce and iterate content yourself	Staff who want the full workflow
Platform contribution	Claude Code + terminal	Suggest improvements, build templates, automate processes	Technically interested staff

Level 1 is a complete, permanent, professional way of working. Not a stepping stone.

The ideal destination for those who want it is Claude Code in VS Code — it gives you the most direct control over the platform. But it takes time to reach, and not everyone needs to get there. Progress at your own pace.

15. Frequently Asked Questions

Will this replace my job? No. The platform replaces formatting, version management, and multi-format distribution — not your expertise, judgment, or relationships. The Centre needs what you know and who you know. It does not need you to spend your week formatting or finding Word documents.

Do I have to learn to code? No. Claude Co-Work is a conversation in plain English. You describe what you need; it handles the rest. The key skill development needed is getting into the mindset of giving conceptually clear instructions, informed by a basic understanding of how large language models work.

Can I still use Word for my own drafts and notes? Personal scratchpads are your business. But project notes, meeting records, and research material should go into the platform. This is the information base your AI worker draws on — the more context it has, the better it works.

What happens to existing Word documents? They stay where they are. The platform is for new work going forward. Existing documents are only migrated if actively needed for a current project.

What if I find a mistake in the output? Request the change — in Claude Co-Work or by telling a colleague. It is fixed and resubmitted for your review. This is a normal part of the workflow, not an emergency.

What if I really struggle with this? That is a legitimate response, not a personal failing. Individual support is available. Raise concerns early — and often.

What are AI workers? Each staff member is complemented by a dedicated collection of carefully defined AI workers that know your projects, can draft content, execute analyses, search the entire information base, and spawn additional copies for parallel tasks. Think of it as a highly capable postdoctoral researcher that has infinite knowledge, but poor contextual understanding and social skills.

Are AI workers really better than specialists? For most analytical tasks, yes. Large language models now surpass the vast majority of technical specialists in producing competent analysis across law, economics, environmental science, and policy. This is the reality of knowledge work in 2026. What they cannot do is apply your judgment, maintain your relationships, or make strategic decisions. Your role moves up a level — you guide, enhance, and verify rather than produce from scratch.

Why do I need a GitHub account? It is how you log in to the internal site and how changes are tracked. You do not need to learn GitHub's interface — think of it as your login credential for the platform.

What is the difference between the internal site, the public site, and the client sites? Same content, different audiences. The internal site shows everything to staff. The public site shows content marked as public — the Centre's external-facing research. Client sites show only that client's project, gated by email address. You control which audience sees what through a single metadata field in the content (`visibility`).

What about SharePoint and UNSW requirements? Final outputs (branded PDFs and DOCX files) can be downloaded from the internal site and copied to SharePoint as needed for long-term record keeping. Automated upload to SharePoint is desirable and technically straightforward, but is currently blocked by UNSW IT restrictions on API access to institutional SharePoint tenancies. This is a roadmap item — when access is granted, the platform can push final outputs to SharePoint automatically. In the interim, the manual step is straightforward: download from the internal site, upload to SharePoint. Use of SharePoint for collaborative drafting or working document storage will be retired, except for the limited context of sensitive human resources matters. The platform replaces SharePoint as the Centre's primary document management and collaboration system.

16. Why This Matters

The substantive capability of large language models has now surpassed the vast majority of technical specialists for most analytical tasks across the disciplines that define the Centre's work — law, economics, environmental science, and policy analysis. Every peer organisation has access to the same models. The models are a commodity. The differentiator is the infrastructure to deploy them systematically and the consolidated information base that gives AI workers full context — every note, transcript, meeting record, and research file.

The FTA Modernisation Project and the Steamships/Swire tender — both delivered through agentic approaches on this platform — achieved approximately 90 per cent operational cost savings compared to traditional production methods. These are not projections; they are measured results from completed work.

What remains distinctly human is precisely what the Centre was established to provide: deep domain expertise, institutional relationships built over years of engagement, strategic judgment about which questions matter and where resources create the most impact, and the credibility that comes from a track record of rigorous independent analysis. The platform does not diminish these capabilities. It frees the people who possess them from spending the bulk of their time on formatting, version management, and document production.

A team of 20+ staff complemented by 20+ dedicated AI workers — each able to spawn copies as needed — is not an incremental improvement. It is a different kind of organisation. Ad hoc use of AI in a browser window is not a strategy. A platform that consolidates the Centre's entire information base and converts AI capability into consistent, branded, access-controlled outputs — deployed to the right audiences automatically — is infrastructure. That is the transition proposed.

17. Productivity, Staffing, and the Case for April

Where capacity is released

The platform automates tasks that currently consume a large share of researchers' time:

- Formatting and layout — fully automated by the build pipeline.
- Multi-format distribution — one approval produces web pages, PDFs, and DOCX simultaneously.
- Access control and infrastructure provisioning — fully automated from content metadata.
- First-draft production — substantially automated with human specification and review.
- Version management — eliminated by the platform's version control.

The capacity released can be directed in several ways: produce more with the same team, maintain current output with fewer resources, or redirect time toward engagement, field research, and the relationship work that defines the Centre's value. These are strategic decisions for leadership.

The AI-native workforce

The shift described in this document has implications for how the Centre is staffed:

- Subject matter specialists remain essential — but the bar rises substantially. The Centre needs fewer researchers who are individually more capable. With AI workers handling production, the premium is on judgment, relationships, and the ability to direct complex work. Staff who can orchestrate multiple AI workers across a project are multiplicatively productive. The Centre needs people who are genuinely good at what they do — and can direct AI workers to amplify that.
- Operations functions that previously required dedicated FTEs are substantially automated. Financial reporting, administrative coordination, compliance reporting — tasks that consumed one or more full-time positions can now be handled by an AI worker under staff direction. The person overseeing finance, for example, shifts from compiling reports to reviewing AI-produced analysis and focusing on strategic financial planning. This is not about eliminating roles — it is about rebalancing so that the team is collectively focusing on higher-value work.
- A full-time junior software developer joins the operations team. This person would handle platform troubleshooting, build issues, first-line support for staff learning the workflow, and process improvement. This is a structural addition — the role is essential for the platform to function smoothly and for staff to be supported properly.
- The communications team shifts to template and format design. Rather than producing individual outputs, the comms team designs the reproducible templates, branding, and formats that the platform applies automatically — ensuring consistency across all outputs while freeing the team from repetitive production work.

Why April, and why all at once

The platform is built and working. It is not a prototype. On deployment it is immediately capable of publishing real research outputs to real audiences with real access control, building on production experience from the FTA Modernisation project, and Solomon Islands Ocean Accounts Project (the latter generating excellent lessons learned about use of Github across technical skill levels).

A clean transition — all new projects go through the platform from a fixed date — is less disorienting than an indefinite ambiguous period of "should I use the old way or the new way?" The first fortnight is difficult. The support described in this document exists specifically because the adjustment is real and is taken seriously.

18. Risks and Limitations

AI quality is managed, not assumed. Current-generation AI workers include self-improvement and adversarial verification layers that catch the vast majority of errors that plagued earlier models — fabricated citations, invented statistics, internal inconsistencies. The platform's multi-agent architecture means outputs are checked before they reach you. This does not eliminate the need for human review — it changes its character. Your review is strategic: does the argument hold, is the framing appropriate for this audience, does the analysis reflect reality as you understand it from the field. You are reviewing for judgment, not proofreading for hallucinations. That said, factual claims in published outputs should still be verified by someone with domain knowledge. The platform makes this faster, not optional.

Vendor dependence exists. The platform relies on Cloudflare (hosting, access control, storage), GitHub (version control), and AI model providers (content generation). These are stable, widely-used services, but the dependence is acknowledged. Core content is stored in plain markdown files in a git repository — the most portable and durable format available.

Not everything needs the platform. Email, Slack messages, and calendar stay where they are. But substantive project work — even informal notes and meeting records — belongs in the platform because it feeds the information base that AI workers draw on. The line is simple: if it contributes to a project, it goes in the platform. If it is a personal reminder or a quick message, it does not.