TheAICommand WC AI Manual

Working with AI under the SRC Act, in seven parts

Version 1.0 15 August 2026

Contents

Before you start

This manual takes a workers compensation practitioner from nothing to competent with AI under the Safety, Rehabilitation and Compensation Act 1988. It is assembled entirely from guidance already published at theaicommand.com, so every statement in it can be traced to a page you can read, and from there to the primary source that page cites. The source index in the back matter names each one.

Nothing here is new research. Where a legal claim is restated, it keeps the wording of the published article it came from. Where a detail changes over time, an indexed rate, a lump-sum cap, a vendor retention term, a file version, this manual points you at the page that carries the current figure rather than printing a number that will be wrong by the time you read it. The volatile-details register in the back matter lists every one of those, so you always know what to check.

The statutory basis

Section references in this manual have been checked against the current compilation of the Act, Compilation No. 82, register C2026C00285, in force from 1 July 2026, as at 15 August 2026. Verify any section you rely on against the current compilation at legislation.gov.au before you act on it.

This manual covers the Commonwealth scheme only. State schemes run on entirely separate legislation, and one of the first failure modes of using a general-purpose model on claims work is that it will blend the two into one confident paragraph. Part 3 exists partly to stop that.

The disclaimer stack

> **Content disclaimer**
> This manual is for general educational purposes only and does not constitute legal advice, liability determination guidance, or a substitute for professional judgement. Workers compensation decisions must be made by appropriately qualified and authorised persons under the Safety, Rehabilitation and Compensation Act 1988. All AI outputs described in this manual require human review before use in any claims management context.

> **Educational tool only**
> Educational tool only. SRC Act section references are general practitioner education, not legal interpretation. Do not enter claimant identifiers at any point.

> **Not an audit instrument**
> This tool is for personal capability reflection only. Do not share the report as evidence in performance discussions or as a compliance audit instrument. Results are not a regulatory finding.

The Verified Draft Method

TheAICommand works to the Verified Draft Method: de-identify the inputs, ground the model in your own source material, keep a person at the decision point, verify against the primary source, and log what happened. The steps run in order and each one earns its place. Skip the first and the rest are unsafe. Skip the last and you cannot prove any of it happened.

1. De-identify the inputs. This is the step that makes every step after it possible. Part 2 of the Method is where Part 4 of this manual begins.
2. Ground the model in your own source material. A model working from its training data will invent provisions. A model working from your verified section summaries will not.
3. Keep a person at the decision point. The practical test is simple: if a regulator or a tribunal would ask who decided, the answer has to be a name, not a tool.
4. Verify against the primary source. Every section reference in every draft is checked against the Act before the draft goes anywhere.
5. Log what happened. Four lines on the file, written every time.

The site publishes a worked example of the Method in this practice: a section 57 medical examination referral, where AI builds the referral and the delegate decides the requirement, and where the requirement has been a reviewable determination since 14 June 2024. Each of the seven parts of this manual names the step of the Method it serves.

The placeholder card

This page is designed to be the one you keep beside you. A shared placeholder convention is what makes de-identified work usable at team scale: square brackets, capital letters, underscores instead of spaces, so any colleague can read the working copy and any output can be re-personalised with a controlled find and replace.

- [CLAIMANT_NAME] for the claimant, everywhere their name appears, including possessives and email addresses.
- [CLAIM_NUMBER] for the claim number and any scheme reference or employee identifier.
- [DATE_OF_BIRTH] for the date of birth, which never has a legitimate reason to enter a drafting tool.
- [EMPLOYER_UNIT] for the agency, branch, team or worksite, wherever the unit is small enough to identify the person.

- [TREATING_PRACTITIONER] for the treating practitioner, with [PRACTICE], [IME_PROVIDER], [REHAB_PROVIDER], [SUPERVISOR] and [WITNESS_1] built on the same pattern as needed.

Extend the set rather than improvising. An ad hoc placeholder like "the doctor" or "Ms X" defeats the find-and-replace step at the end and breeds inconsistency across a team. The canonical form for the treating practitioner is [TREATING_PRACTITIONER], used consistently across the articles, the learning module and the prompt pack, because a treating practitioner is often a specialist, physiotherapist or psychologist rather than a general practitioner.

De-identification is mandatory, not optional

No real names, claim numbers or dates of birth ever enter an AI tool. Not in a prompt, not in a pasted document, not in a file upload, not in an enterprise workspace with training excluded. Replace identifiers with placeholders before any material goes near a model, keep the placeholders through every AI step, and re-personalise only inside your controlled claims system after human review. If material cannot be de-identified, it does not go in. There is no exception for urgency, seniority or tool vendor.

Part 1

The landscape

What the Comcare scheme is, where AI has actually landed in claims work, and the one rule that governs everything after this.

What this part covers

- The Commonwealth scheme in one page, and why the state schemes matter to a model
- The five claims workflows where AI is already in use
- The honest counterweight: where the scheme-level promise breaks
- The dividing line this whole manual holds

This part sets up the Method. It decides nothing on its own.

The Comcare scheme in one page

The SRC Act 1988 (Cth) runs the Commonwealth workers compensation scheme. It covers employees of the Australian Government and its agencies, the ACT public sector, and the staff of corporations that hold a self-insurance licence under the Act. Comcare administers claims for many Commonwealth agencies; licensed self-insurers determine their own claims under the same legislation. Determinations are made by authorised decision makers, called delegates, applying the Act to the evidence on file. A claimant who disagrees can seek reconsideration, and then merits review in the Administrative Review Tribunal. Every workflow in this manual is built to survive that review chain.

State schemes run on entirely separate legislation. The first question on any claim is which scheme applies, and the first rule of using AI on claims is that a model trained on the general internet will happily blend state and Commonwealth concepts into one confident paragraph. Your project space, built in Part 3, exists partly to stop that.

Where AI has actually landed

Across the claims lifecycle, AI shows up in five places. None of them is decision making, and the value in every one of them is the same: converting an unstructured pile into structure a person can work with quickly.

1. Intake and initial triage. Reading incoming claim documents, extracting the structured fields a claims system needs, and flagging what is missing before a file is allocated.
2. Decision support on liability. Organising evidence against the statutory tests so a delegate can see what bears on each limb, and what does not yet exist.
3. Communications drafting. Turning a completed decision into plain-English correspondence that a claimant can act on.
4. Document analysis and case file review. Summarising long reports, building chronologies, and mapping where accounts diverge.
5. Quality assurance and pattern detection. Sampling de-identified closed files for consistency in reasoning, structure and correspondence.

Notice what these have in common. Every one of them is collation, structuring or drafting. Collation is where the hours go, which is why AI is genuinely useful here. None of them answers a statutory question.

The honest counterweight

The most attractive AI proposition in workers compensation is also the most dangerous one: predictive triage, where a model scores incoming claims for likely complexity, duration or cost, and the scheme routes resources accordingly. The attraction is obvious. The problem is that a risk score attached to a person, generated before any human has read the file, sits uncomfortably close to the procedural fairness obligations that govern how a claim is handled and how a decision is explained. A claimant is entitled to a decision made on the evidence in their file, not on what a model expects people like them to do.

That does not make predictive work impossible. It makes the boundary a governance question rather than a technical one, and it is the reason Part 6 exists as its own part rather than as a closing paragraph.

The dividing line

One sentence governs this manual and every workflow in it. AI organises the evidence and frames the questions. The delegate answers them.

What AI must never do is answer a statutory question. Whether employment contributed to a significant degree is a weighing of medical evidence. Whether an action was reasonable and reasonably taken is a contextual judgement about what the employer knew and did. Whether liability is accepted is the statutory decision itself. A model that appears to answer these questions produces a confident conclusion with no accountable reasoning behind it, and a delegate who adopts that conclusion has not made the decision the Act requires them to make.

> **Read this part with the rule already in place**
> Nothing in Parts 3 to 5 is safe without Part 2. If you are skimming, skim Part 2 first: it is the one part with no optional content in it.

Part 2

The rules that bind you

The liability architecture, the drafting rule that sits above all others, the twenty-four provisions this practice runs on, and the clocks.

What this part covers

- Sections 5A, 5B and 14, and the three-limb reasoning order
- Why determinations state outcomes and never recommend them
- A reference card for the twenty-four provisions you will meet
- The review chain and the clocks that run on it

This part is the Method step 4: verify against the primary source.

The liability architecture: sections 5A, 5B and 14

Section 5A defines injury. It covers a physical or mental injury arising out of, or in the course of, the employee's employment, a disease, and an aggravation of an injury or disease. The definition then takes away with one hand what it gave with the other: under section 5A(1), a condition is not an injury if it was suffered as a result of reasonable administrative action taken in a reasonable manner in respect of the employee's employment. That carve-out is the most contested territory in the scheme, because it frequently decides psychological claims arising from performance management.

Section 5A(2) gives the carve-out its content. It lists, without limiting the concept, the actions that count as administrative action: a reasonable appraisal of the employee's performance; reasonable counselling action, formal or informal; reasonable suspension action; reasonable disciplinary action, formal or informal; anything reasonable done in connection with any of those actions; and anything reasonable done in connection with the employee's failure to obtain a promotion, reclassification, transfer or benefit, or to retain a benefit. Two distinct reasonableness questions attach to each action: was the decision to take it reasonable, and was it taken in a reasonable manner. Both are evaluative judgements on the evidence, and both belong to the delegate alone.

Section 5B defines disease. A disease is an ailment, or an aggravation of an ailment, that was contributed to, to a significant degree, by the employee's employment. Section 5B(3) sets the bar: significant degree means a degree that is substantially more than material. Most psychological conditions travel through this disease pathway, which means the employment contribution question is answered before the section 5A carve-out is ever reached.

Section 14 is the liability provision. It makes the relevant authority liable to pay compensation in respect of an injury suffered by an employee if the injury results in death, incapacity for work, or impairment. Section 14 is where liability lives, and only there.

> **The most common drafting error in this practice**
> It is a common drafting error, in human and AI drafts alike, to attribute the reasonable administrative action carve-out to section 14. It sits in section 5A(1). A determination that cites the wrong provision for the operative exclusion invites reconsideration on the spot.

The three-limb reasoning order

On the page, the Act reads as definitions followed by a liability rule. At the desk, a case manager works through three limbs in a fixed order, and the order is what keeps the reasoning defensible.

1. Limb 1: Is there a compensable condition? Characterise what is claimed. A physical or mental injury under section 5A, or a disease under section 5B? If disease, does the evidence show employment contributed to a significant degree, substantially more than material? This limb is medical-evidence heavy and turns on what the reports actually say.
2. Limb 2: Is the employment connection made out? For an injury, did it arise out of or in the course of employment? For a disease, the contribution question from limb 1 carries the connection. This limb turns on the factual matrix: what happened, when, where, and what records it.
3. Limb 3: Does an exclusion apply? Chiefly, for psychological claims, was the condition suffered as a result of reasonable administrative action taken in a reasonable manner under section 5A(1), tested against the section 5A(2) categories and both reasonableness limbs. Only if the condition survives all three limbs does section 14 liability attach.

The three-limb order is also the map of where AI helps. Every limb rests on evidence that arrives as an unstructured pile: certificates, reports, emails, file notes, statements. AI is genuinely good at converting that pile into structure that matches the legal test: a document-by-document evidence map, a dated chronology, a table of employer actions tagged against the section 5A(2) categories, a list of gaps and conflicts. That work is collation, and collation is where the hours go.

Determinations state outcomes

One drafting rule sits above all others in this practice, and it needs to be installed before any AI drafting begins, because language models default to hedged, advisory prose. Determinations are decisions, not recommendations. They state outcomes in deterministic language: what is decided, under which provision, on what evidence. They never say what the decision maker thinks the outcome ought to be.

> **The deterministic language rule**
> A determination states its outcome: "Liability is rejected. The evidence does not satisfy section 14 of the SRC Act on the balance of probabilities." Or: "Liability is accepted. The condition is an injury under section 5A to which section 14 applies." The word "should" never appears in outcome wording, because the delegate is not recommending a decision to someone else. The delegate is making it. Every prompt chain in Part 3 carries this rule, and the delegate checks it again before anything issues.

Two qualifications on that rule, because over-applying it causes its own errors. First, it binds specimen determination text, not ordinary advisory prose: this manual says "you should run the sweep on Monday" without difficulty. Second, section 37 uses the word in its own statutory language, allowing the authority to determine that an employee should undertake a rehabilitation program, in the statute's own words. When you quote section 37, the qualifier travels with the quote.

The reference card

The twenty-four provisions below are the ones this practice runs on. They are reproduced here so the manual works away from a screen, and so the file library in Part 3 has a starting point. They are educational summaries, not the Act.

Educational summaries of the provisions this manual works in. Section 40 is covered in Part 4 rather than here.

Section	What it does
s4	Interpretation. The definitions section. Defines suitable employment, impairment, permanent, ailment and aggravation, and routes injury to section 5A and disease to section 5B.
s5	Employees. Who the Act covers: Commonwealth and Commonwealth authority employees, employees of licensed corporations, and deemed categories such as declared volunteers.
s5A	Definition of injury. An injury is a disease (which takes its own test from s5B), a physical or mental injury arising out of or in the course of employment, or an aggravation of such an injury. The closing words of s5A(1) exclude conditions resulting from reasonable administrative action taken in a reasonable manner; s5A(2) gives non-exhaustive examples of such action.
s5B	Definition of disease. An ailment, or an aggravation of an ailment, contributed to, to a significant degree, by the employment. Significant degree means substantially more than material; s5B(2) lists matters that may be taken into account.
s6	Injury arising out of or in the course of employment. Deems injuries sustained in listed circumstances (ordinary recess, directed travel, approved study, attending treatment or examinations) to be employment connected. Ordinary home-to-work commuting is excluded.
s7	Provisions relating to diseases. Deeming and presumption rules for diseases, including the date a disease injury is sustained, firefighter cancer presumptions and first responder PTSD presumptions.
s8	Normal weekly earnings. How NWE before injury is calculated from average weekly hours, ordinary time rates and allowances over the relevant period, with rules for regular overtime, part-time work, increments and indexation.
s14	Compensation for injuries. The core liability provision. Comcare is liable to pay compensation for an injury that results in death, incapacity for work, or impairment. Exclusions apply for intentionally self-inflicted injury, and for serious and wilful misconduct unless the injury results in death or serious and permanent impairment.
s16	Compensation in respect of medical expenses. Liability for the cost of medical treatment it was reasonable for the employee to obtain in the circumstances, of an amount Comcare determines appropriate, plus journey and accommodation rules.
s19	Compensation for incapacity. The main weekly payments provision. NWE minus ability to earn during maximum rate weeks, stepping down to adjustment percentages between 75 and 100 per cent after the 45-week equivalent cap.
s20	Incapacity with a superannuation pension. Weekly compensation for a retired employee on a superannuation pension: the notional s19 amount reduced by the employer-funded superannuation amount and 5 per cent of NWE.
s21	Incapacity with a lump sum benefit. The lump sum counterpart of s20: the notional s19 amount reduced by weekly interest on the lump sum and 5 per cent of NWE.

Section	What it does
s21A	Incapacity with a pension and a lump sum. Where the retired employee received both a superannuation pension and a lump sum, both offsets apply against the notional s19 amount.
s22	Employee maintained in a hospital. A substituted weekly rate where an employee with no dependants has been maintained in a hospital or similar place for at least a year, set between half and the full otherwise payable amount.
s23	When incapacity payments are not payable. Weekly payments stop at pension age (with a 104-week rule where the injury is suffered at or after the age 2 years before pension age), during imprisonment on conviction, and on redemption under s30.
s24	Permanent impairment. Lump sum compensation for permanent impairment, assessed under the approved Guide and expressed as a percentage of the maximum amount. A 10 per cent threshold applies (5 per cent for binaural hearing loss), with finger, toe, taste and smell exceptions.
s28	Approved Guide. Comcare prepares, and the Minister approves, the Guide to the Assessment of the Degree of Permanent Impairment. The Guide binds Comcare, licensees and the Administrative Review Tribunal.
s36	Rehabilitation capability assessment. The rehabilitation authority may, and on the employee's written request must, arrange an assessment of capability of undertaking a rehabilitation program. Refusing or failing, without reasonable excuse, to undergo an examination suspends compensation rights, other than s16 medical treatment.
s37	Rehabilitation programs. The rehabilitation authority may determine that an employee should undertake a rehabilitation program, weighing the listed factors. Refusing or failing, without reasonable excuse, to undertake the program suspends rights (other than s16 medical treatment) until it begins.
s53	Notice of injury. The Act does not apply to an injury unless written notice is given as soon as practicable. Savings apply where the authority is not prejudiced, or the failure arose from death, absence from Australia, ignorance, mistake or another reasonable cause. Claims are made separately under s54.
s57	Medical examinations. The relevant authority may require an examination by one nominated medical practitioner, complying with the approved assessments and examinations Guide. Refusing or failing, without reasonable excuse, or obstructing the examination suspends rights.
s62	Reconsideration of determinations. The internal review step. A determining authority may reconsider a determination on its own motion, or must on a request made within 30 days of the determination coming to notice (extendable). The outcome is a reviewable decision.
s63	Notice of reviewable decisions. A reviewable decision must be served in writing with its terms, its reasons, and a statement of the right to apply to the Administrative Review Tribunal.
s64	Applications to the ART. A reviewable decision may go to the Administrative Review Tribunal; a primary determination must pass through s62 reconsideration first. The 60-day lodgement window sits in s65(4).

> **Reference card**
>
> Educational reference only. Not legal advice. Always consult the current authoritative version of the SRC Act.

The review chain and its clocks

Workers compensation review under the SRC Act sits in three tiers, and the order matters. The first review of any primary determination is internal reconsideration under section 62, carried out by a reconsideration officer who was not involved in making the primary determination. That is a procedural fairness requirement, not a courtesy. The reconsideration officer can affirm, vary or revoke the primary determination, and the reconsideration produces a reviewable decision, which is the artefact that triggers the next tier.

A reviewable decision can then be reviewed on the merits by the Administrative Review Tribunal, which replaced the AAT on 14 October 2024. The Tribunal can affirm, vary, set aside and substitute its own decision, or set aside and remit the matter, with some limits on substitution for certain rehabilitation decisions. From there, an appeal lies to the Federal Court on a question of law only, which is narrower than it sounds: errors in applying a legal test, misconstructions of a provision, denials of natural justice and inadequate reasons can all ground a question of law, but disagreements about the weight given to evidence generally cannot.

Three clocks run on that chain, as at 15 August 2026. A request for reconsideration of a primary determination should generally be made within 30 days of the determination being notified. An application to the ART for review of a reviewable decision must generally be made within 60 days of the reviewable decision being received. An appeal to the Federal Court must be lodged within 28 days of the ART decision. Extensions are possible in narrow circumstances and should never be assumed.

Separately, the prescribed periods for determining a claim have run since 1 April 2024: 20 calendar days for an injury claim and 60 for a disease claim, with 30 days for a reconsideration. AI can run those clocks and flag what is approaching. It cannot decide whether an extension is warranted, and it cannot decide what a claim is.

> **Why this matters for AI work**
> Every clock in this section is a date arithmetic problem, which is exactly the kind of work a model does reliably and a human does inconsistently at volume. It is also the kind of work where a wrong answer is invisible until it is expensive. Run the sweep, then check the dates against the file.

Part 3

The tools

Why a project space beats an ad hoc chat, the locked instruction text, the four files that ground the model, and what never goes in.

What this part covers

- The case for a configured project space over a chat window
- The custom instructions, reproduced in full and treated as locked
- The four-file reference library and what goes in each
- Skill files, and the material that never enters a tool at all

This part is the Method step 2: ground the model in your own source material.

Why a project space beats ad hoc chats

A loose chat window is the wrong tool for determination support, for three reasons. It carries no standing rules, so the de-identification gate and the never-decide boundary have to be retyped, and eventually are not. It carries no reference material, so the model free-associates about workers compensation from its training data, blending state schemes, superseded provisions and other countries' law into plausible prose. And it leaves no consistent structure for a team to review, so every AI-assisted determination looks different at audit.

A project space fixes all three at once: one instruction set, four reference files, four chains. Every conversation starts inside the guardrails. The site publishes a ready-to-paste configuration for this exact set-up, including the platform-specific settings and the verified vendor facts, at theaicommand.com/library/configurations. Vendor retention and training terms change; that page carries the current verified extract, which is why this manual does not print them.

The custom instructions

The instructions are the contract. They encode the de-identification gate, the never-decide boundary, the section-referencing discipline and the deterministic language rule, so no individual prompt has to remember them. Paste this text as the project instructions and treat it as locked: changes go through whoever owns your team's AI governance, not through individual edits mid-claim.

The de-identification gate prompt that enforces rule 1 is the first prompt in Part 5. Install it at the top of every claims conversation until the habit is automatic.

The locked instruction text itself, six hard rules, is reproduced in the LM-W01 learning module and in the platform configuration on the site. Its six rules are: all inputs are de-identified and placeholders are preserved; the assistant drafts and never decides; every legislative reference names a specific section and no case law is ever cited; draft determination text uses deterministic language with outcomes stated and never recommended; no diagnosis, capacity or medical conclusion is ever inferred, with gaps recorded rather than filled; and every draft ends with a "For human review:" list of the highest-risk points the delegate must verify against the claim file.

The file library

Four small markdown files give the model verified reference material and your team's actual formats. Keep them short, current, and owned: each file has one person responsible for updates, and every entry in the key-sections file is verified against the current compilation on the Federal Register of Legislation before it goes in.

- src-act-key-sections.md: the sections your team works with daily, each with a verified plain-language summary and the test it sets. This is the file that stops the model inventing or blending provisions. The reference card in Part 2 is your starting point.
- determination-templates.md: your team's determination structures for section 14, 16, 19 and 24 decisions, including the standard review rights wording, so drafts land in the format reviewers expect.
- evidence-summary-template.md: the evidence map and chronology table formats from Part 4, so every claim's working documents look the same at audit.
- de-identification-rules.md: the placeholder set, the five-stage workflow and the excluded-material list, so the tool can enforce the gate from its own reference file.

Every entry in the key-sections file carries an owner and a last-verified date. That is not bureaucracy: it is the only thing that stops a superseded summary sitting in a file the model treats as authoritative for two years.

Skill files: the next step up

A prompt is something you paste. A skill file is something you install: a named, versioned instruction file with a fixed six-part anatomy, purpose, when to use, inputs required, method, output format and guardrails, that the tool loads every time it does that task. The advantage over re-prompting is consistency. The same input produces the same shape of output for every person on the team, and the guardrails are built in rather than remembered.

Two skill files for this practice are published and free to download at theaicommand.com/library/skills: a de-identified claim chronology builder, and a determination evidence check that audits a draft's evidence chain before a decision is made. The library page carries the current version and the honest tested-against record for each, which is why this manual does not print either.

What never enters an AI tool, even de-identified

De-identification removes direct identifiers. It does not make every document safe, because some material identifies a person through the shape of its facts rather than the name on its cover, and some material carries legal sensitivities that no redaction cures. Three categories stay out of AI tools entirely.

- Psychiatric report narratives with unique identifying fact patterns. A psychiatric history often reads like a fingerprint: a specific sequence of workplaces, family events, treatments and incidents that could identify the person to anyone who knows them, with every name removed. Summarise what you need by hand; the narrative itself does not travel.

- Third-party medical opinions under dispute. Where an IME opinion or a treating practitioner report is actively contested, the document is potential evidence in a review. Feeding it through an external tool creates handling questions you do not want to answer at the Administrative Review Tribunal. Work from your own hand-written summary of the competing positions instead.
- Anything subject to a legal hold. Material preserved for litigation, a review, or an investigation is under a preservation obligation. It is not working material, and it does not enter any tool outside the controlled environment, de-identified or not.

> **AI tools**
> AI tool recommendations are illustrative and not endorsements. Your organisation must apply its own AI governance framework. De-identification, human-in-the-loop review, and APP compliance are non-negotiable. The tool does not replace professional judgment.

Part 4

The workflows

Nine desk workflows in one fixed shape: what triggers it, what goes in, the steps, what the AI produces, what stays human, and the file note.

What this part covers

- De-identify, build the chronology, draft a section 14 determination
- Map administrative action, map conflicting medical opinions
- Cross-check an incapacity calculation, assemble an NWE evidence pack
- Build a section 57 referral, plan a return to work under sections 36, 37 and 40

This part runs the whole Method, once per workflow.

Each workflow below follows the same six-line shape. Read the shape once and the nine become quick to scan: trigger, inputs, steps, what the AI produces, what stays human, and the file note that closes it. The file note is the same four lines every time, so it is stated once in Part 6 and referenced here.

Workflow 1. De-identify

Trigger: any claim material about to leave your controlled system. Inputs: the source document, untouched. This workflow runs before every other workflow in this manual, without exception.

The de-identification workflow has five stages, and the order is the control. Each stage has one job, and the real identifiers exist only at the two ends, inside controlled systems.

1. Source document. The original stays untouched in the case management system. It is never edited, never pasted anywhere, never attached to anything outside the controlled environment.
2. De-identified working copy. Save a clearly marked copy, run a find-and-replace sweep for names, numbers, dates of birth, units and practitioners, then visually scan paragraph by paragraph for leakage: signature blocks, letterheads, email subject lines, file names, distinctive phrases. Only the clean working copy travels.
3. AI-assisted draft. The working copy enters the approved tool. All prompting, structuring and drafting happens in placeholder form. If a real identifier surfaces mid-conversation, stop, delete the conversation, and start again from a clean copy.
4. Human review. The delegate reviews and edits the output with placeholders intact. Discussion with colleagues happens in placeholder form too, which keeps every intermediate artefact privacy-safe.
5. Re-personalisation, in controlled systems only. The final, human-approved text is re-identified with a controlled find and replace inside the case management system. This is the last step, never an intermediate one, and it never happens inside the AI tool.

The sweep itself clears five identifier categories, and they are deliberately broader than the strict legal definition of personal information because the goal is robustness rather than minimum compliance: direct claimant identifiers; indirect identifiers such as a job title small enough to single someone out, a site name or a combination of very specific dates; treating practitioner identifiers including practice name and suburb; third-party identifiers such as witnesses, family members, IME providers and lawyers; and free-text leakage, which is the trickiest because it hides in plain sight in signature blocks, letterhead text and file metadata.

What a clean working copy looks like

A fictional incapacity summary, every identifier replaced and every fact preserved: claimant [CLAIMANT_NAME], claim [CLAIM_NUMBER]; lumbar strain, liability accepted under section 14 on [DETERMINATION_DATE]; [TREATING_PRACTITIONER] certifies capacity for suitable duties, 4 hours per day, 3 days per week, no lifting above [LIFT_LIMIT], review in 4 weeks; graduated return at [EMPLOYER_UNIT] since [RTW_START_DATE], currently 12 hours per week against a pre-injury 38. The placeholders keep the document fully workable. The structure, the dates pattern, the restrictions and the section references all survive, and nothing in it could identify a person.

Stays human: the visual scan. A find-and-replace sweep catches what you told it to catch. Only a person reading paragraph by paragraph catches the distinctive phrase that identifies someone to a colleague who knows the team.

Workflow 2. Build the chronology

Trigger: a file with more than a handful of documents, before any drafting starts. Inputs: the de-identified working copies, each with a stable document label.

1. Read every supplied document and extract each event that carries a date or a sequence position.
2. Order the events using the date tokens supplied, never by inference about what must have happened first.
3. Attribute every entry to the document that records it, by label, so each line can be walked back.
4. Flag every gap, every undated event and every point where two documents describe the same event differently.
5. Return a neutral table. No narrative, no characterisation, no conclusion about what the sequence shows.

What the AI produces: a dated table with a source column and a flagged-gaps section. What stays human: everything the table is for. The chronology is an index for your own reading of the file, never a substitute for it, and the gaps it flags are prompts for your enquiry rather than findings.

Why the chronology is the right first workflow

It is the highest-volume, lowest-judgement task in the practice, which makes it the safest place to calibrate your trust in a tool. Run it on a closed file first and compare the output against how the claim was actually worked.

Workflow 3. Draft a section 14 determination

Trigger: the evidence is assembled and you are ready to reason. Inputs: the de-identified evidence map and chronology from workflow 2, inside the project space from Part 3.

1. Sort the evidence against the three limbs. The prompt for this is the first determination prompt in Part 5.
2. Walk the section 5A(2) categories where a psychological claim is in issue, as questions for you to answer rather than answers.
3. Draft the legislative framework section from your verified section summaries, never from the model's own knowledge.
4. Produce a determination skeleton with every finding left as a bracketed space.
5. Complete every bracketed finding yourself, in your own words and your own reasoning.
6. Run the deterministic-language and identifier check over the completed draft.
7. Read every source document the draft relies on before the determination issues.

What the AI produces: structure, framework text and a skeleton. What stays human: every finding, every outcome, every weighing of evidence, and the decision itself. The skeleton is deliberately full of holes, because the holes are the job.

Workflow 4. Map administrative action under section 5A

Trigger: a psychological claim where employer action is in issue. Inputs: the de-identified employer action record and the claimant's account.

1. Build an inventory of every employer action in the period, with the date and the document recording it.
2. Tag each action against the section 5A(2) categories: appraisal, counselling, suspension, disciplinary action, anything done in connection with those, or anything done in connection with a failure to obtain or retain a benefit.
3. For each action, list the evidence bearing on whether the decision to take it was reasonable, and separately the evidence bearing on whether it was taken in a reasonable manner.
4. List what is missing under each action before a delegate could form a view.

What the AI produces: a two-column evidence inventory against the statutory categories. What stays human: both reasonableness questions, on every action, without exception. These are contextual judgements about what the employer knew and did, and a model has no access to that context.

Workflow 5. Map conflicting medical opinions

Trigger: two or more medical opinions that do not agree. Inputs: your own de-identified summaries of each opinion, not the reports themselves where any of them is under dispute.

1. Set the opinions side by side on the questions actually in issue, one row per question.
2. Separate what each practitioner observed from what each practitioner concluded.
3. Mark where the opinions agree, where they diverge, and where one is silent.
4. List the questions a supplementary request would need to ask to resolve each divergence.

What the AI produces: a comparison table and a question list. What stays human: the weighing. Which opinion is preferred, and why, is the reasoning a reconsideration officer and then a Tribunal will read. A model that offers a preference is generating a medical judgement without a clinician.

Workflow 6. Cross-check an incapacity calculation

Trigger: a completed section 19 incapacity calculation, before payment. Inputs: your calculation, plus the de-identified inputs it was built from.

The principle here is the reverse of every other workflow: you calculate first, then the AI recomputes independently and the two are compared. Asking a model to produce the figure invites you to check its work, which is the wrong way round. Asking it to check yours keeps the calculation, and the accountability, with you.

1. Do the calculation yourself, in full, and write down the inputs you used.
2. Give the model the same inputs and the method, and ask it to compute independently.
3. Compare the two figures. Where they differ, find out why before changing anything.
4. Ask the model to list the assumptions it made that you did not state.
5. Record the check, and the resolution of any difference, on the file.

What the cross-check catches: arithmetic slips, a step applied in the wrong order, an input carried forward incorrectly. What it cannot catch: an input that is wrong on the file. If the normal weekly earnings figure was wrong going in, both calculations will be wrong coming out, and they will agree with each other.

Workflow 7. Assemble a normal weekly earnings evidence pack

Trigger: a normal weekly earnings figure needs to be determined or revisited. Inputs: de-identified payroll records, employment records and any variation history.

1. List every source document that bears on earnings, with its date range and what it evidences.

2. Set out the earnings pattern the documents show, period by period, with the source for each.
3. Flag every discontinuity: a change in hours, a change in classification, an absence, an allowance that starts or stops.
4. List what is missing before the figure can be settled.

What the AI produces: an evidence pack, which is an organised set of what the documents say. What stays human: the figure. The pack exists so a delegate can determine normal weekly earnings on the evidence, and the single most important design rule for these prompts is that the model is never asked what the figure is.

Workflow 8. Build a section 57 examination referral

Trigger: an examination is being considered. Inputs: the de-identified issues in dispute and the evidence gap the examination would fill.

1. State the issues actually in dispute, from the file, before drafting a single question.
2. Draft questions that ask for observation and opinion within the practitioner's expertise.
3. Run a legal-conclusion detector over the draft questions and remove any that ask the practitioner to answer the statutory test.
4. Assemble the documents list, with a relevance note for each, so the practitioner knows why each is there.

What stays human: whether to require the examination at all. Since 14 June 2024 that requirement has been a reviewable determination, which means the decision to make it is itself subject to the review chain in Part 2. Build the referral with AI; make the requirement yourself, and record the reasons.

Workflow 9. Plan a return to work under sections 36, 37 and 40

Trigger: an accepted claim with incapacity or impairment. Inputs: the de-identified certificate, the role demands, and the barriers the employee has described.

Section 36 allows the rehabilitation authority to arrange an assessment of an employee's capability of undertaking a rehabilitation program where an injury results in incapacity for work or impairment. Section 37 allows the authority to determine that an employee should undertake a rehabilitation program, in the statute's own words, and sets out the matters to consider, including any written assessment, cost, the likely improvement in employment opportunity, the likely psychological effect on the employee of not providing the program, the employee's attitude to the program, and alternative programs. Section 40 places the duty on the employer's side: where an employee is undertaking, or has completed, a rehabilitation program, the employer takes all reasonable steps to provide suitable employment or to assist the employee to find it, with suitable

employment defined in section 4 by reference to the employee's age, experience, training, language and other skills, and their suitability for rehabilitation or vocational retraining.

Notice what kind of decisions these are. The section 36 assessment is a clinical and vocational evaluation. The section 37 determination weighs statutory considerations, including a psychological one. The section 40 duty is a judgement about what is reasonable for a particular employer and a particular employee. Every one of them is a human decision resting on professional assessment. AI's role in this part is narrower than in workflow 3, and the boundary is brighter: barrier analyses, plan options, case conference agendas and progress summaries are structured documents built from information others have provided, and that is the whole of it.

> **Three hard edges in return to work planning**
>
> Never let AI infer diagnosis or capacity: the certificate says what it says, and a model that extrapolates from a lifting restriction to a view about what the employee can really manage is generating a medical opinion without a clinician. Never generate medical conclusions: summarising what a report states is support work, but concluding what the reports mean clinically, reconciling conflicting certificates, or predicting recovery timeframes is medical judgement. Never draft section 37 program requirements without the rehabilitation authority's own assessment: a program drafted from a model's assumptions rather than an actual assessment is indefensible at review.

Psychological claims sharpen every edge. Section 37 expressly includes the likely psychological effect on the employee of not providing a program among its considerations, and planning for psychological injury turns on trust, pacing and consultation. AI must never be used to characterise an employee's language as resistance, reluctance or non-compliance. Draft neutral questions, quote certificates, and leave the reading of people to the people in the room.

Part 5

The prompts

Twelve prompts reproduced verbatim from the published pack, covering all six prompt types and all four practitioner tiers.

What this part covers

- How to read a prompt: setup, the prompt, what to never do, what good looks like, how to use it
- Twelve prompts spanning determination, evidence, file review, communication, examination briefing and behavioural engagement
- The de-identification gate first, always
- How to build your own library from these twelve

This part serves Method steps 1 and 3: de-identify, and keep a person at the decision point.

How to read a prompt

Every prompt below carries five parts, and all five matter. The setup states what must be true before you paste anything, and it always includes de-identification. The prompt itself is copy-paste ready, with merge fields in square brackets. What to never do is the boundary: the thing that would turn support work into a decision. What good output looks like is a rubric of three to five criteria you check the output against, which is what stops you accepting a fluent answer that is wrong. How to use it is the step sequence, and it always ends at a human review gate.

These twelve are reproduced verbatim from the prompt pack published on this site. Nothing has been shortened, and no rubric bullet has been dropped to make a page fit: a truncated rubric would be a worse prompt, and it would make the verbatim claim false. Where a prompt runs past a page break, it continues on the next page rather than being cut.

The twelve are chosen to a stated rule so that you can extend the set the same way. Every prompt type appears at least once. Every practitioner tier appears at least once. The de-identification gate is always first. Each prompt pairs to a workflow in Part 4. And every one is available to the four case-manager paths, so the set is not leader-only.

Building your own library

Save these into the project space from Part 3 and refine the wording as your team's formats evolve. A prompt library is a living artefact, and the discipline that keeps it useful is the same as the file library: one owner, a version, and a rule that a prompt earns its permanent slot by being used rather than by being written.

Two rules on claim data, both of which the prompts below already encode. First, prompts are written and tested against fictional scenarios, never against a live file, so a library can be shared and reviewed without any privacy question attaching to it. Second, every prompt carries its merge fields in the same bracketed form as the placeholder card, so a prompt and a de-identified working copy fit together without editing either.

> **Prompt pack**
> Apply your organisation's AI governance framework. Never paste claimant identifiers into a public AI tool. Human-in-the-loop review is mandatory.

The twelve

De-identification verification gate before any pasting

Tier: Learning File review

Pairs with workflow 1. Run this before anything else, every time.

Setup

Replace names, claim numbers, dates of birth, addresses and employer identifiers with placeholders before pasting anything; this prompt is the gate that verifies you actually did. Run it at the start of any AI session that will touch claim material, before a single document is pasted. It costs a minute and prevents the worst failure mode.

The prompt

```
Before I paste any claim material into this conversation, run my de-identification
checklist. Ask me to confirm each line, one at a time, and do not proceed until all
six are confirmed.

1. Names: claimant, family members, witnesses, supervisors and treating
practitioners are all replaced ([CLAIMANT_NAME], [TREATING_PRACTITIONER],
[SUPERVISOR], [WITNESS_1]).
2. Numbers: claim number, employee ID and payroll number are all replaced
([CLAIM_NUMBER], [EMPLOYEE_ID]).
3. Dates: date of birth is replaced ([DATE_OF_BIRTH]). Injury and treatment dates
remain only where they cannot identify the person in combination with other details;
otherwise [INJURY_DATE].
4. Places: employer unit, worksite and practice names are replaced ([EMPLOYER_UNIT],
[WORKSITE], [PRACTICE]).
5. Leakage: email signatures, letterheads, file names, subject lines and document
properties are stripped.
6. The excluded list: the material contains no psychiatric report narrative with a
unique identifying fact pattern, no disputed third-party medical opinion, and
nothing subject to a legal hold. If any of these are present, stop: they do not
enter this tool in any form.

If at any point in this conversation you see what looks like a real name, claim
number or date of birth, stop immediately and tell me to de-identify before doing
anything else.
```

What to never do

Never skip the gate because the session feels routine; identifiers leak through signatures, file names and metadata more often than through body text. If any line cannot be confirmed, nothing is pasted.

What good output looks like

- The checklist runs one line at a time and waits for confirmation
- Nothing is pasted before all six lines are confirmed
- The standing stop instruction is acknowledged for the whole session
- Excluded-list material is kept out entirely, not de-identified harder

How to use it

1. Open every claim-material AI session with this prompt. 2. Confirm each line honestly against the material you are about to paste. 3. If any line fails, stop and fix the material first. 4. Human review gate: the gate

checks your preparation; you remain responsible for every paste that follows, and for stopping the session if an identifier slips through.

Three-limb evidence sort for a section 14 determination

Tier: Learning Determination support

Pairs with workflow 3, step 1. The three-limb sort is where determination work starts.

Setup

Replace names, claim numbers, dates of birth, addresses and employer identifiers with placeholders before pasting anything. Prepare a de-identified summary of the claim evidence: the condition claimed, the key documents, and the employment circumstances asserted. This prompt organises material only; it never produces a view on liability.

The prompt

```
I am a workers compensation delegate working under the Safety, Rehabilitation and
Compensation Act 1988. I will paste a de-identified summary of claim evidence for a
[CLAIM_TYPE] claim. Sort it against the three-limb liability structure for a section
14 determination. Do not weigh anything and do not suggest an outcome.

1. Limb 1 (compensable condition): list every piece of evidence bearing on what the
condition is and, if a disease, on whether employment contributed to a significant
degree under section 5B. Present competing evidence side by side as neutral
summaries.
2. Limb 2 (employment connection): list the evidence on how the condition connects
to employment: events, dates, duties, locations, and the document recording each.
3. Limb 3 (exclusions): list anything that could engage section 14(2) or 14(3), and
any employer action that could fall within section 5A(2), with the evidence
describing each.
4. Gaps: list what is missing under each limb before a delegate could form a view.

Output four labelled lists. No conclusions, no recommendations.

[DE_IDENTIFIED_SUMMARY]
```

What to never do

Never paste claimant names, claim numbers or dates of birth. Never treat the sorted lists as findings; the delegate assesses the evidence and decides, and nothing goes to file without your review.

What good output looks like

- Every listed item names the source document it came from
- Competing evidence appears side by side with no weighing language
- Section references are limited to sections 5A, 5B and 14
- The gaps list is present and names specific missing material
- No outcome, recommendation or credibility comment appears anywhere

How to use it

1. Run your de-identification checklist over the summary before pasting. 2. Paste the prompt, then the summary. 3. Check every item in the output against the claim file and correct anything misattributed. 4.

Delete any weighing or outcome language that crept in. 5. Human review gate: use the verified lists as working material for your own assessment only; nothing enters the file until you have reviewed and adopted it as your own work.

Chronology build with gap flags

Tier: Learning Evidence organisation

Pairs with workflow 2. Gaps are flagged, never filled.

Setup

Replace names, claim numbers, dates of birth, addresses and employer identifiers with placeholders before pasting anything; keep only the dates that cannot identify the person in combination with other details, otherwise use placeholders like [INJURY_DATE]. Prepare the same de-identified file digest used for the evidence map. The chronology records events; it draws no conclusions.

The prompt

```
From the de-identified claim material I paste, build a chronology for a matter under the Safety, Rehabilitation and Compensation Act 1988.

1. Events: every event in date order, one line each, with the document that records it named alongside.
2. Gap flags: flag any unexplained gap in the timeline, stating the period and what is silent about it.
3. Uncorroborated events: flag any event asserted with no contemporaneous record, naming the single source it rests on.
4. Date conflicts: flag any event where documents give different dates, quoting both.

Record events neutrally: what happened according to which document, never why or with what significance. Do not weigh the evidence. Do not suggest an outcome.

[DE_IDENTIFIED_FILE_MATERIAL]
```

What to never do

Never rely on the chronology without checking each entry against its source document; a misplaced date changes the shape of a claim. Never paste claimant identifiers or date combinations that could identify the person.

What good output looks like

- Every event names the document recording it
- Gaps, single-source events and date conflicts are each flagged separately
- Entries are neutral descriptions with no significance commentary
- Conflicting dates are quoted, not resolved

How to use it

1. De-identify the material, replacing identifying dates with placeholders. 2. Paste and generate the chronology. 3. Verify every entry and flag against the source documents. 4. Human review gate: resolve date conflicts from the file yourself; the chronology is preparation for your assessment, never evidence in itself.

Plain-English determination letter

Tier: Learning Claimant communication

Pairs with workflow 3, after the decision is made. Plain English is a comprehension task, not a softening one.

Setup

Replace names, claim numbers, dates of birth, addresses and employer identifiers with placeholders before pasting anything. Have the final de-identified determination text ready, placeholders intact, in ChatGPT Enterprise, Copilot, Claude, or an equivalent approved tool.

The prompt

```
I am a workers compensation delegate working under the Safety, Rehabilitation and
Compensation Act 1988. I have made a determination and I will paste the
de-identified text with placeholders intact. Draft a plain-English letter to
[CLAIMANT_NAME] explaining it.

1. State what was decided. The outcome, reasons and section references must match
the determination exactly. Change the language, never the meaning.
2. Explain the key reasons in short sentences. Any scheme term that must stay gets a
one-line explanation.
3. State what happens next and any action required, with the key date unmissable.
4. Copy the reconsideration and review rights from the determination accurately. Do
not paraphrase timeframes. If the determination text contains no rights statement,
or the 30-day timeframe is missing, stop and flag that as a defect for human review
rather than drafting around it.
5. End with "For human review:" listing every sentence where plain English risked
changing the legal meaning.

Constraints: keep the deterministic framing. The letter explains what has been
decided and why. Do not speculate about other outcomes, do not soften the decision
into a suggestion, and do not add anything the determination does not contain.

[PASTE_DEID_TEXT]
```

What to never do

Never paste identifiers; placeholders only. The letter never leaves your hands unreviewed: verify every section reference and timeframe against the determination and the Act before sending.

What good output looks like

- Outcome, reasons and section references match the determination exactly
- Sentences are short and every retained scheme term is explained
- Reconsideration and review rights are copied accurately, not paraphrased
- One clear next step and key date stand out
- The "For human review:" list flags genuine meaning-shift risks

How to use it

1. De-identify the determination text against your checklist. 2. Paste the prompt, then the text. 3. Compare the draft letter to the determination line by line for meaning changes. 4. Verify timeframes and section

references against the Act. 5. Human review gate: read the letter as the claimant would, correct anything unclear, and only then move it into your correspondence process.

MI call plan for an RTW check-in

Tier: Learning Behavioural engagement

Pairs with workflow 9. Rehearsal, never a script to read.

> **Behavioural ethics note**
> Behavioural techniques support the claimant's own decisions and recovery. They must never be used to engineer a choice, manufacture agreement, or advance scheme outcomes at the claimant's expense.

Setup

Replace names, claim numbers, dates of birth, addresses and employer identifiers with placeholders before pasting anything. Write a short de-identified context block: where the return to work is up to, what was agreed last time, and the purpose of this check-in, in ChatGPT Enterprise, Copilot, Claude, or an equivalent approved tool.

The prompt

```
I am preparing for a return to work check-in call with [CLAIMANT_NAME]. From the
de-identified context below, draft a motivational interviewing informed call plan.

1. Opening: two open questions that invite [CLAIMANT_NAME]'s own account of how
things are going, not yes/no questions.
2. Reflections: three example reflective listening responses I could adapt, each
feeding back meaning rather than parroting words.
3. Information exchange: an ask-offer-ask structure for anything I need to share:
ask what they already know, offer the information with permission, ask what they
make of it.
4. Closing: a summary structure that collects what was said and confirms any next
step in [CLAIMANT_NAME]'s words where possible.
5. Watch-outs: three moments where I might slip into the fixing reflex, arguing for
change instead of evoking their reasons, with an alternative move for each.

Constraints: the plan supports [CLAIMANT_NAME]'s own decisions about recovery. Flag
any wording that pressures, guilt-frames, or removes choice. This is preparation,
not a script.
```

What to never do

Never paste identifiers. The plan informs a human conversation: you deliver the call, respond to what is actually said, and never read the plan at the claimant. Never act on it unreviewed.

What good output looks like

- Openings are genuinely open questions
- Reflections feed back meaning, not parroted words
- Information exchange follows ask-offer-ask
- Fixing-reflex watch-outs come with alternative moves

- No line pressures, guilt-frames or removes choice

How to use it

1. Write the de-identified context block. 2. Run the prompt. 3. Adapt the plan into your own words; keep only what sounds like you. 4. Human review gate: internalise it, put it aside, and hold the conversation as a person, following up on what the claimant actually says.

Section 14 determination skeleton with bracketed findings

Tier: Applying Determination support

Pairs with workflow 3, step 4. Every finding is left as a bracketed space on purpose.

Setup

Replace names, claim numbers, dates of birth, addresses and employer identifiers with placeholders before pasting anything. Have your reviewed, de-identified evidence map and chronology ready, plus your standard determination template wording. The skeleton carries structure only; every finding and the outcome stay blank for the delegate.

The prompt

```
Using the approved, de-identified evidence map and chronology I paste, draft a
skeleton for a section 14 determination under the Safety, Rehabilitation and
Compensation Act 1988 with this structure:

1. Decision heading. Leave the outcome line blank for me.
2. Claim background: [CLAIMANT_NAME], claim [CLAIM_NUMBER], the condition claimed,
the date of claim.
3. Legislative framework: the sections in play ([SECTION_REFS]), summarised from my
verified section summaries only.
4. Evidence considered: the documents from the evidence map, listed neutrally.
5. Findings structure: a heading for each limb (condition, employment connection or
contribution, exclusions), the relevant evidence slotted under each, and a bracketed
space for my finding: [DELEGATE_FINDING].
6. Determination: [OUTCOME_AS_DETERMINED_BY_DELEGATE], followed by the notice
content required by section 61: the terms, the reasons, and the right to request
reconsideration under section 62.

Use deterministic language in every sentence that could survive to final text. Do
not complete any [DELEGATE_FINDING] or the outcome. End with "For human review:" and
the three highest-risk points I must verify against the claim file.

[VERIFIED_SECTION_SUMMARIES]

[EVIDENCE_MAP_AND_CHRONOLOGY]
```

What to never do

Never let the skeleton state or imply an outcome; every finding and the outcome are the delegate's alone. Never paste claimant identifiers, and never issue anything built on the skeleton without full review against the claim file.

What good output looks like

- Every finding space and the outcome remain bracketed and blank

- Evidence sits under the limb it bears on, with source documents named
- Legislative statements match the verified summaries and name their sections
- The section 61 notice elements are present
- The "For human review:" list names concrete verification points

How to use it

1. Confirm the evidence map and chronology were reviewed and approved by you first. 2. Paste the prompt with the map, chronology and section summaries. 3. Check every evidence slot against the claim file. 4. Complete every finding and the outcome yourself, in your own words. 5. Human review gate: verify the section references and notice content, then read the full determination end to end before it goes anywhere.

Conflicting medical opinions side by side, no weighing

Tier: Applying Evidence organisation

Pairs with workflow 5. Side by side, and no weighing.

Setup

Replace names, claim numbers, dates of birth, addresses and employer identifiers with placeholders before pasting anything, including all practitioner names ([TREATING_PRACTITIONER], [IME_PROVIDER], [SPECIALIST_1]). Prepare de-identified extracts of each conflicting opinion. The table structures the conflict; resolving it is the delegate's assessment.

The prompt

```
I will paste de-identified extracts of medical opinions that appear to conflict, for
a matter under the Safety, Rehabilitation and Compensation Act 1988. Build a
side-by-side structure. Do not weigh the evidence, do not prefer any opinion and do
not suggest an outcome.

1. Issue rows: identify each issue the opinions address (for example diagnosis,
cause, contribution of employment under section 5B, capacity).
2. For each issue, set the opinions side by side in columns: what each says, quoted
or closely paraphrased, with author role and document date.
3. Stated basis: under each opinion, what the author says it rests on: examination,
history taken, documents reviewed, assumptions made.
4. Common ground: note, neutrally, any point on which the opinions agree.
5. Delegate questions: list the questions I need to resolve about each conflict,
including whether any opinion rests on facts not in evidence.

Output the table and lists. No preference language anywhere.

[DE_IDENTIFIED_OPINION_EXTRACTS]
```

What to never do

Never let the structure hint at which opinion is stronger; resolution is the delegate's reasoning task on the whole file. Never paste claimant or practitioner identifiers.

What good output looks like

- Each conflict is broken down by issue, not by document
- Every opinion carries its stated basis

- Common ground is recorded neutrally
- No preference, weight or credibility language appears
- Delegate questions target the basis of each opinion

How to use it

1. De-identify all opinion extracts. 2. Paste and generate the side-by-side structure. 3. Verify every quoted position against its source report. 4. Human review gate: write your own reasoning on each conflict from the full reports; the table maps the disagreement, it never resolves it.

Question-set review against the claim's issues

Tier: Applying Examination briefing

Pairs with workflow 8. Questions that ask for a legal conclusion are the failure mode to catch.

Setup

Replace names, claim numbers, dates of birth, addresses and employer identifiers with placeholders before pasting anything. Have two things ready: the draft referral questions, and a de-identified list of the issues actually open on the claim.

The prompt

```
I will paste a draft set of examination referral questions, then a de-identified
list of the issues open on the claim. Review the questions against the issues.

1. Coverage map: a table matching each open issue to the questions that address it.
Flag every issue with no question, and every question serving no open issue.
2. Legal-conclusion check: flag any question asking the examiner to decide
liability, whether the statutory employment connection or contribution test is met,
or whether an administrative action was reasonable. Those belong to the delegate.
3. Answerability: flag questions the examiner cannot answer from examination and the
enclosed documents, and say what is missing.
4. Precision: flag compound questions to split and vague questions to sharpen, with
suggested rewording.
5. Sequence: propose the order the questions run best in.

Constraints: do not add new issues beyond my list. Suggested rewording keeps every
question clinical. Where a question is fine, say so and move on.

[PASTE_DRAFT_QUESTIONS]
[PASTE_DEID_OPEN_ISSUES_LIST]
```

What to never do

Never paste identifiers. The issues list drives the whole review, so an inaccurate list produces a confident but wrong coverage map; the delegate owns the final question set. Never send unreviewed.

What good output looks like

- Every open issue and every question appears in the coverage map
- Legal-conclusion questions are caught and reworked as clinical
- Unanswerable questions are flagged with what is missing
- Reworded questions are sharper and still clinical

How to use it

1. Write the de-identified open-issues list from the file. 2. Run the review. 3. Decide each flag yourself; the coverage map advises, it does not decide. 4. Human review gate: finalise the question set against the file before the referral is issued.

Reasons quality pass: findings-to-evidence mapping

Tier: Mastering Determination support

Pairs with workflow 3, step 6. This is Method step 4 in prompt form.

Setup

Replace names, claim numbers, dates of birth, addresses and employer identifiers with placeholders before pasting anything. Copy your full de-identified draft reasons, placeholders intact. This pass checks the linkage between findings and evidence; it never touches the merits.

The prompt

```
I will paste my de-identified draft reasons for a determination under the Safety,
Rehabilitation and Compensation Act 1988. Run a quality pass on the reasoning
structure only. Do not comment on the merits and do not redraft anything.

1. Findings inventory: quote each finding sentence exactly as drafted.
2. Evidence linkage: under each finding, list the evidence the draft ties to it.
Flag any finding with no cited evidence.
3. Orphaned evidence: list evidence discussed in the draft but tied to no finding.
4. Leaps: flag, as questions for me, any point where the draft moves from evidence
to conclusion without stating the connecting reasoning.
5. Outcome wording: flag any outcome sentence that is not stated in the present
tense or that recommends rather than states.
6. Section references: list every section cited so I can verify each against the
current compilation.

Output six labelled parts and nothing else.

[DE_IDENTIFIED_DRAFT_REASONS]
```

What to never do

Never act on a flag without checking it against the claim file; the pass finds structural gaps, the delegate judges whether they matter. Never paste claimant identifiers, and never adopt suggested framings because none are given: the fixes are drafted by you.

What good output looks like

- Every finding is quoted exactly, not paraphrased
- Unsupported findings and orphaned evidence are both surfaced
- Leaps appear as questions, not corrections
- No merits commentary or redrafting appears
- The section list matches a manual read of the draft

How to use it

1. De-identify and paste the full draft reasons. 2. Work each flag against the claim file. 3. Redraft weak

linkages in your own words. 4. Verify every listed section against the current compilation. 5. Human review gate: reread the revised reasons end to end; the quality pass supplements your review, it never replaces it.

Reconsideration outcome letter in plain English

Tier: Mastering Claimant communication

Pairs with the review chain in Part 2. A reconsideration outcome carries its own review rights.

Setup

Replace names, claim numbers, dates of birth, addresses and employer identifiers with placeholders before pasting anything. Have the de-identified reconsideration decision ready, including whether the determination was affirmed, revoked or varied under section 62(5).

The prompt

```
I will paste a de-identified reconsideration decision made under section 62 of the
SRC Act. Draft the letter that communicates it to [CLAIMANT_NAME].

1. State plainly whether the original determination was affirmed, revoked or varied,
and what that means in practical terms for the claimant.
2. Explain the reasons, matching the decision exactly. Where the reconsideration
reached a different view from the original, explain the difference without
criticising the original decision-maker.
3. Where the determination was varied, set out precisely what changed and what
stayed the same.
4. State the next review right accurately: this decision is a reviewable decision,
and an application to the Administrative Review Tribunal may be made within 60 days
after being served with this notice, under section 65(4), with extensions available
under the Tribunal's rules.
5. End with "For human review:" listing sentences where the affirmed, revoked or
varied distinction could be misread.

Constraints: deterministic framing; the letter states what has been decided. No
speculation about how the Tribunal might view the matter.

[PASTE_DEID_COMPLAINT_SUMMARY]
```

What to never do

Never paste identifiers. The affirmed, revoked or varied distinction and the 60-day ART timeframe must be exact; a wrong timeframe here can cost a claimant their review. Never send unreviewed.

What good output looks like

- The section 62(5) outcome is stated plainly with its practical effect
- Variations are itemised: what changed, what stayed
- The ART right and 60-day timeframe are stated accurately
- No speculation about Tribunal outcomes appears

How to use it

1. De-identify the reconsideration decision. 2. Run the prompt. 3. Verify the outcome characterisation and

the 60-day timeframe against the decision and the Act. 4. Human review gate: final letter checked through your reconsideration correspondence process before sending.

Reconsideration readiness check

Tier: Mastering File review

Pairs with Part 2. Run it before a file goes to reconsideration, not after.

Setup

Replace names, claim numbers, dates of birth, addresses and employer identifiers with placeholders before pasting anything. Prepare a de-identified summary of the reconsideration request, the original determination and the file state. The check confirms the process is set up properly; it says nothing about the outcome.

The prompt

```
A reconsideration under section 62 of the Safety, Rehabilitation and Compensation
Act 1988 is about to proceed. From the de-identified summary I paste, run a
readiness check as questions. No view on any ground and no view on the outcome.

1. Request validity: does the summary show a written request setting out the reasons
for it, and was it given within 30 days after the determination first came to the
notice of the person requesting, or has the determining authority allowed a further
period (section 62(3))? Frame what the summary shows and what needs confirming.
2. Independence: is the reconsideration allocated to a person who was not involved
in making the original determination, as section 62(4) requires?
3. Material assembly: is everything that was before the original decision maker
assembled, and is new material lodged with the request logged separately?
4. Notice trail: is the original notice under section 61 on file, showing the terms,
reasons and reconsideration right as issued?
5. Outstanding items: a consolidated list of what to confirm or obtain before the
reconsideration starts.

Output five labelled parts.

[DE_IDENTIFIED_RECONSIDERATION_SUMMARY]
```

What to never do

Never treat the readiness check as clearance on the timing question; whether the request was in time, or a further period is allowed, is decided by the determining authority on the facts. Never paste claimant identifiers.

What good output looks like

- The timing check states what the summary shows without deciding it
- The section 62(4) independence requirement is checked explicitly
- Original and new material are confirmed as separated
- No view on any ground appears

How to use it

1. Build the de-identified summary from the file. 2. Paste and run the check. 3. Verify every item against the file, especially the timing facts. 4. Resolve outstanding items before allocation. 5. Human review gate: the

reconsideration delegate confirms independence and material completeness personally before starting the fresh review.

Peer review checklist and coaching questions for a draft determination

Tier: Mentoring Determination support

Pairs with Part 6. Peer review is a governance control, not a favour.

Setup

Replace names, claim numbers, dates of birth, addresses and employer identifiers with placeholders before pasting anything, with the drafting colleague's agreement. Copy the colleague's de-identified draft determination. The output is a question set for a peer conversation, never a corrected draft.

The prompt

```
I am peer reviewing a colleague's de-identified draft determination under the
Safety, Rehabilitation and Compensation Act 1988. Build me a structured review
checklist and a coaching question set from the draft I paste. Do not redraft
anything and do not give a view on the merits.

1. Structure check: does the draft move from the question each section asks, to
findings, to supporting evidence, to the conclusion the section compels. Note where
it departs, as observations.
2. Evidence linkage: list any finding without cited evidence and any evidence tied
to no finding.
3. Outcome wording: flag any sentence that recommends rather than states an outcome.
4. Section references: list every section cited, for the author to verify against
the current compilation.
5. Procedural fairness markers: note whether the draft records that adverse material
was put to the claimant and the response considered, as questions.
6. Coaching questions: convert every observation into an open question I can ask the
author, framed so the author reaches the fix themselves.

Output six labelled parts.

[COLLEAGUES_DE_IDENTIFIED_DRAFT]
```

What to never do

Never run a colleague's draft through the tool without their agreement, and never paste claimant identifiers. The question set supports a conversation; the author owns every change, and no output is kept on the claim file.

What good output looks like

- Every observation is converted into an open coaching question
- No redrafted text or merits view appears
- Evidence linkage gaps are specific, quoting the draft
- Procedural fairness markers appear as questions

How to use it

1. Get the author's agreement and confirm the draft is de-identified. 2. Paste and generate the checklist and questions. 3. Verify each observation against the draft yourself. 4. Hold the conversation using the questions; the author decides every change. 5. Human review gate: the author reviews and reworks their own draft; nothing from the tool is filed.

Part 6

The governance controls

The controls that turn a fast workflow into a defensible one: the human-in-the-loop gate, the file note, disclosure at the Tribunal, and what a team lead owns.

What this part covers

- The three commitments that make an AI-assisted determination defensible
- The four-line file note, written every time
- AI disclosure in expert evidence at the Administrative Review Tribunal
- What a governance lead owns, and the ten-item practice checklist

This part is the Method steps 3 and 5: the person at the decision point, and the log.

The human-in-the-loop gate

The project space produces working drafts faster than any process it replaces, which is exactly why the gate at the end matters more, not less. Three commitments make an AI-assisted determination defensible, and all three are the delegate's personal work.

1. The delegate reads every source document. The evidence map is an index for that reading, never a substitute for it.
2. The delegate treats AI output as a working draft only. Every factual assertion is walked back to the document that records it, every section reference is checked against the Act, and every finding space is completed in the delegate's own reasoning and words.
3. The statutory decision is the delegate's. The outcome, the findings and the reasons belong to the person the Act makes accountable, who signs knowing the reasoning will be read at reconsideration and, if it comes to it, at the Administrative Review Tribunal.

The practical test for all three is the one the Method states: if a regulator or a tribunal would ask who decided, the answer has to be a name, not a tool.

The file note

> **A file note closes the loop**
> Record the AI assistance on the file in four lines: the tool used, the narrow purpose it served, confirmation that inputs were de-identified, and a statement that the delegate reviewed and edited the output so the issued text reflects the delegate's own reasoning. Four lines, written every time, is the difference between a workflow you can defend at audit and one you have to explain from memory.

The note is short by design. A long note invites editing and eventually stops being written. Four lines takes under a minute, and it is the only artefact that survives a change of case manager, a change of system, and the two years between a determination and a Tribunal hearing.

AI disclosure at the Tribunal

The Administrative Review Tribunal now requires expert reports to declare generative AI content. That obligation lands on this practice in three distinct places, and the third one is the one most often missed.

- At commissioning. Build the disclosure requirement into the referral itself, so the practitioner knows the expectation before writing rather than after. This is the cheapest place to solve it, and it is workflow 8, step 4.

- On a report you already hold. A report that carries no disclosure is not automatically unusable, but the absence is a fact about the evidence, and it is a fact the delegate should know before relying on it.
- On your own side of the file. Your AI use in preparing the material that goes to the Tribunal is subject to the same transparency expectations. The four-line file note is what makes that answerable.

What a governance lead owns

Individual discipline is necessary and not sufficient. A handful of controls sit above the desk and belong to whoever owns AI governance for the team.

- A register of the AI uses actually in play, kept current rather than written once, so the answer to "what are we using this for" is a document and not a conversation.
- Sample review of AI-assisted determinations, on the same footing as any other quality sample, looking at reasoning and evidence chains rather than at whether AI was used.
- Ownership of the four reference files, with a named owner and a last-verified date on each, because a superseded section summary in a grounded file is worse than no file at all.
- A privacy assessment covering the specific tools in use, which is the document that answers the question this manual cannot: whether a given tool is approved in your environment.
- A briefing for delegates that states the never-decide boundary in the organisation's own words, so the boundary is a local rule rather than a book they read.

> **Governance**
> Identified risks are illustrative. Apply your own professional judgment and consult your organisation's governance framework.

The claims-practice checklist

Run this before any AI-assisted document leaves your desk

- [ ] No real names, claim numbers or dates of birth ever enter an AI tool. Placeholders only, every time.
- [ ] The excluded list is respected: unique psychiatric narratives, disputed third-party opinions and legal hold material never travel, even de-identified.
- [ ] Every claims conversation starts inside the project space, under the locked instructions, never in an ad hoc chat.

- [ ] Every section reference in every AI draft is verified against the Act before the draft goes anywhere.
- [ ] AI output is structured as questions and drafts. Any output that states or implies an outcome is discarded, and the prompt is tightened.
- [ ] The delegate reads every source document. The evidence map is an index, not a substitute.
- [ ] Determination text uses deterministic language: outcomes are stated, never recommended, and "should" never appears in outcome wording.
- [ ] Every finding and every outcome is completed by the delegate, in the delegate's own words.
- [ ] Return to work drafting quotes certificates exactly and never infers diagnosis, capacity or motivation.
- [ ] A four-line file note records tool, purpose, de-identification and human review on every AI-assisted document.

Part 7

The ninety-day plan

What to do on Monday, what to build in the first month, what to run on live work after that, and what to measure at ninety days.

What this part covers

- Days 1 to 7: five actions that take under an hour in total
- Days 8 to 30: build the project space and the file library
- Days 31 to 60: move onto live work with the checks running
- Days 61 to 90: measure, and decide what earns a permanent slot

This part sequences the whole Method into a working habit.

Days 1 to 7: do this Monday

1. Print the placeholder card from the front of this manual, tape it to your monitor, and run the five-stage de-identification workflow on one real document as a practice pass: source, working copy, sweep, visual scan, and stop there. The habit forms at the desk, not in a policy document.
2. Create the determination-drafting project space: one project in your approved tool, the locked custom instructions pasted in, and empty stubs for the four reference files. Twenty minutes, once.
3. Draft src-act-key-sections.md for the five sections your team touches most, verifying every summary against the current compilation on the Federal Register of Legislation before it enters the file.
4. Run the chronology workflow on one closed, de-identified claim file and compare the output against how the claim was actually worked. A closed file is a risk-free test bed, and the gaps it exposes calibrate your trust in the tool.
5. Write the four-line AI file note for that practice run, and agree the note format with your team lead so every AI-assisted document from here carries one.

Days 8 to 30: build the ground

The first month is construction, and all of it happens on closed files or fiction. Nothing in this window touches a live determination, which is what makes it safe to be wrong.

- Finish the file library. Four files, each with an owner and a last-verified date. The section summaries are the load-bearing one: the reference card in Part 2 is your starting point, but every entry is verified against the current compilation before it goes in.
- Install the two published skill files and run each on a closed file. A skill file that produces a shape your team does not use is a skill file to adapt, not to abandon.
- Work through the twelve prompts in Part 5 in order, on fiction. The point is not to use all twelve; it is to find the three you will actually use weekly.
- Run the de-identification gate prompt at the top of every one of those sessions, so the habit is automatic before it matters.

Days 31 to 60: move onto live work

The second month puts the workflows on live files, with every check running. Two things change: the volume goes up, and the cost of a missed control goes from zero to real.

- Adopt one standing routine. A weekly deadline sweep against the prescribed periods and the review clocks in Part 2 is the highest-value candidate, because it is pure date arithmetic and the failure mode is silent.
- Run the determination workflow end to end on a live file, with the evidence check before the decision and the deterministic-language pass after it.
- Write the four-line file note every time, without exception. If the note is being skipped, the workflow is moving faster than the governance and something needs to slow down.
- Take one AI-assisted determination to peer review using the twelfth prompt in Part 5, and treat the reviewer's questions as the calibration they are.

Days 61 to 90: measure and prune

The third month is where a workflow either earns its place or is dropped. Measure honestly, because the failure mode here is a tool that feels fast and produces work that takes longer to check than it took to write.

- Which workflows saved time, and which added a checking burden larger than the drafting they replaced. Keep the first group; drop the second without sentiment.
- Whether reasoning quality moved. The evidence check and peer review from month two are your evidence base here, not your impression.
- Whether the file note is being written every time. If it is not, the control has failed, and the fix is process rather than reminder.
- What your section summaries got wrong. Every correction you made to a model output because the file said something different is a gap in the grounding, and closing it is the cheapest improvement available.

At ninety days, the practical measure of success is narrow and worth stating plainly: the same decisions, reached the same way, with the collation done faster and the reasoning trail easier to follow at reconsideration. Nothing in this manual is aimed at changing what a delegate decides.

> **Where to go next**
> The WC Self-Assessment at theaicommand.com/tools/wc-self-assessment gives you a role-specific reading of where your practice sits and a personalised prompt pack. Learning module LM-W01 covers this ground as an assessed course with a thirty-question final. The workflow recipes, skill files and platform configurations are all free at theaicommand.com.

Source index

Every part of this manual is assembled from the pages below, all published at theaicommand.com and all live at the time of writing. Read any of them to check a statement in this manual, and follow its citations to check the statement against the primary source.

Method and reference

- /method: the Verified Draft Method, the five steps and one worked example per profession.
- /glossary: the Australian reference layer, including the SRC Act provision pages that carry the current compilation identity. Sections 8, 24 and 29 each have their own page.
- /library/skills: the versioned skill file library, including the claim chronology builder and the determination evidence check.
- /library/configurations: the platform configurations, including the workers compensation Claude Projects set-up and its verified vendor facts.
- /recipes: the workflow recipe library, including determination drafting with de-identification controls and the Monday claims deadline sweep.
- /learning-hub/modules/lm-w01-ai-for-src-act-claims-practice: LM-W01, the assessed course this manual expands.
- /tools/wc-self-assessment: the self-assessment that generates a role-specific report and prompt pack.

Workers compensation articles

- /workers-comp/src-act-and-ai-assisted-determinations: the five-step frame, file note conventions, and the current compilation identity.
- /workers-comp/de-identification-toolkit-for-case-managers: the five identifier categories, the placeholder convention and the daily desk routine.
- /workers-comp/ai-tools-in-claims: the five claims workflows where AI is already in use.
- /workers-comp/art-review-rights: the three review tiers and their clocks.
- /workers-comp/src-act-statutory-timeframes-and-ai: the prescribed periods and where AI helps run them.
- /workers-comp/ai-and-the-reasonable-administrative-action-exclusion: section 5A, both reasonableness limbs, and what never to automate.
- /workers-comp/predictive-analytics-and-claims-triage: the honest counterweight on scheme-level prediction.

- /workers-comp/claim-chronology-skill-file: the chronology method, output format and guardrails.
- /workers-comp/offline-wc-evidence-chronology-builder: neutral chronology design and the field patterns to avoid.
- /workers-comp/determination-evidence-check-skill-file: auditing a draft determination's evidence chain before the decision.
- /workers-comp/ai-mapping-conflicting-medical-opinions: setting opinions side by side without weighing them.
- /workers-comp/treating-practitioner-reports-and-ai: summarise reports, never weigh them.
- /workers-comp/incapacity-cross-check-workflow: the five-step cross-check and what it cannot detect.
- /workers-comp/normal-weekly-earnings-ai-evidence-pack: what belongs in an evidence pack and how to stop AI choosing the answer.
- /workers-comp/section-57-medical-examination-referrals-with-ai: the referral workflow and the reviewable requirement.
- /workers-comp/ai-section-16-reasonable-medical-treatment: building the section 16 reasonableness picture without making the call.
- /workers-comp/src-act-section-36-ai-rehabilitation-assessment: the section 36 workflow, step by step.
- /workers-comp/ai-recovery-at-work-and-suitable-duties: consultation as a control, not a courtesy.
- /workers-comp/ai-drafted-motivational-interviewing-scripts: rehearsal aids and the ethics boundary.
- /workers-comp/art-expert-evidence-ai-disclosure: the disclosure requirement and where it lands.
- /workers-comp/wc-prompt-libraries-and-human-review: building a prompt library on fiction, and the review gates.
- /workers-comp/ai-permanent-impairment-evidence-s24: assembling section 24 evidence against the approved Guide.
- /workers-comp/src-act-section-29-ai-household-services-evidence-map: the section 29 evidence map.
- /workers-comp/reasonable-excuse-under-the-src-act-and-ai: the reasonable excuse test across five sections.
- /workers-comp/preventing-double-payment-under-the-src-act-and-ai: third-party and state scheme overlap.

- /workers-comp/ai-drafted-determinations-reasoning-trail-case-note: what a broken reasoning trail costs, and the review clocks.
- /workers-comp/src-act-review-2025-and-ai: the reform horizon and how to triage the recommendations.

The volatile-details register

> **Currency**
> Information current as of generation date. Always verify against the authoritative source.

These are the details that change. This manual deliberately does not print the ones marked "not printed here", because a number that is wrong is worse than a pointer that is right. Check each against the page named before you rely on it.

- Statutory rates and lump-sum caps, including the section 17 death benefit, the section 24 permanent impairment maximum and the section 29 household services and attendant care caps: not printed here. Section 13 of the Act indexes them. Check Comcare's current statutory rates before quoting any figure. The provision pages at /glossary/src-act-s24 and /glossary/src-act-s29 explain which amounts section 13 reaches.
- Indexation dates for weekly compensation and superannuation pensions: not printed here, for the same reason.
- Vendor retention and training terms for any AI platform: not printed here. The current verified extract is on the platform configuration page at /library/configurations.
- Skill file versions and their tested-against records: not printed here. The library page at /library/skills carries both.
- Recipe and configuration tested-on dates: not printed here. Each page carries its own.
- The SRC Act compilation identity: Compilation No. 82, register C2026C00285, in force 1 July 2026, as at 15 August 2026. Verify at legislation.gov.au.
- The prescribed periods, 20 calendar days for injury and 60 for disease, with 30 days for reconsideration, in force since 1 April 2024: as at 15 August 2026.
- The review clocks, 30 days to reconsideration, 60 days to the Tribunal and 28 days to the Federal Court: as at 15 August 2026.
- The ART expert evidence disclosure requirement, commenced 2 March 2026: as at 15 August 2026.
- The section 57 requirement as a reviewable determination, since 14 June 2024: as at 15 August 2026.

- The SRC Act Review, 124 recommendations, delivered to government 25 September 2025 and released 12 December 2025, with no government response and no bill: as at 6 July 2026, which is earlier than the rest of this register. Check before relying on it.

About this manual

TheAICommand WC AI Manual, version 1.0, 15 August 2026. Every part of this manual is assembled from guidance already published at theaicommand.com. The source index in the previous part names each piece so you can check any statement against the page it came from, and against the primary source that page cites.

The current version of this manual, and any superseded versions, are at theaicommand.com/library/manuals.

TheAICommand. Intelligence, At Your Command.