

Workflows, References, AI Protocols, RHE Standards.

The full briefing for NSWRA and RHE assessor leadership — what the tool does at every stage, every document it cites, every AI guardrail, every assurance posture.

drfa-assess.com · Feb 2026 · RHE Project Consulting

Four pillars. Twenty-two slides.

PILLAR 1 • WORKFLOWS

- ▶ 03 The operating principle

- ▶ 04 EW Finalisation

- ▶ 05 IRW Scope Approval

- ▶ 06 IRW Finalisation

- ▶ 07 EPAR Scope Approval

- ▶ 08 EPAR Finalisation

- ▶ 09 The four assistant tools

PILLARS 2 • 3 • 4

- ▶ 10 Referenced documents — EPAR set

- ▶ 11 Referenced documents — Project mgmt, on-costs & redamage

- ▶ 12 Referenced documents — Cross-portal policy & EW/IRW

- ▶ 13 AI protocols — Scope & citation rules

- ▶ 14 AI protocols — Read-only posture & vision audit

- ▶ 15 AI protocols — Precedent corpus

- ▶ 16 AI protocols — Phase trajectory

- ▶ 17 RHE standards — Governance & sign-off

- ▶ 18 RHE standards — Security & access

- ▶ 19 RHE standards — Audit trail & data residency

- ▶ 20 RHE standards — Deployment & operations

- ▶ 21 Closing — what the assessor signs

- ▶ 22 Where to go next

The tool sorts. The tool surfaces options. The assessor — every time — decides.

WHAT RAAGENT DOES

- ▶ Ingest the ZIP / Drive link, classify every file, route by portal type

- ▶ Parse the SmartyGrants assessment artefact (EW Spreadsheet, EPAR SQP Record, Form 306 legacy)

- ▶ Walk evidence folders; build a damage / completion / pre-disaster photo registry per asset

- ▶ Reconcile a 15,000-row labour ledger in 1.8 s against the on-cost fact sheet

- ▶ Cross-check rates against Cost Estimation Tool v2.1

- ▶ Surface ambiguities to a `needs_review` queue

- ▶ Generate the canonical NSWRA artefact at sign-off

WHAT THE SQP / ASSESSOR DOES

- ▶ Sort `needs_review` queue (portal type, missing fields, asset-evidence mismatches)

- ▶ Pick one of three decision options per **Accept · RFI · Reject** line:

- ▶ Confirm or override every CET rate decision

- ▶ Tick D1–D6 + C1–C4 (Decision Confirmation) and S1–S6 (SQP Certification)

- ▶ Sign the EW Assessment Spreadsheet or EPAR SQP Assessment Record

THE LINE THAT DOES NOT MOVE

The AI never marks an outcome. The assessor clicks. The click is what is written to the assessment record — and what is written to the audit trail.

EW Emergency Works — Finalisation

Finalise Cat B Emergency Works claims against the EW Eligibility Criteria — reconcile the ledger, audit defects, sign off the EW Assessment Spreadsheet.

INPUTS THE ASSESSOR RECEIVES

- SmartyGrants EW package (ZIP or Drive link)
- Council labour ledger (.xlsx, up to ~15k rows)
- Defect photos: damage + completion, organised by defect
- Council EW Assessment Spreadsheet draft (or Form 306 legacy)

WHAT RAAGENT AUTOMATES

- Streaming ledger parse — 15k rows in 1.8 s
- On-cost reconciliation vs RA-approved 33.08% (FY24-25) / 34.62% (FY25-26)
- Per-defect photo bucketing + EXIF / GPS / quality check
- Aggregate CET parsing (LGA + project tagging for parent folders)
- Live ingest progress with milestone polling

DECISION OPTIONS SURFACED (EACH LINE)

● **Accept** · ● **RFI** · ● **Reject** · plus per-defect outcome (Compliant · Clarification required · Not eligible) and recommendation.

OUTPUT THE SQP SIGNS

EW Assessment Spreadsheet — 28-column SmartyGrants template, embedded dropdowns, colour-coded Outcome cells, summary block, assessor sign-off.

```
POST /api/submissions/<id>/ew-assessment-form
```

IRW Immediate Reconstruction Works — Scope Approval

Approve the scope of IRW against pre-condition certification and the Standard Treatment Guide v2.1, before delivery begins.

INPUTS THE ASSESSOR RECEIVES

- SmartyGrants IRW Scope package
- Pre-condition certification per asset
- Scope and treatment proposed per asset (mapped to STG v2.1 work types)
- Photographic evidence of damage

WHAT RAAGENT AUTOMATES

- Per-DM bucketing of damage photos with EXIF / GPS / quality checks
- STG v2.1 work-type matching against proposed scope
- Rate cross-check via the Standard Treatment Tool v2.1 (IRW-only)
- Pre-condition completeness flag per asset

DECISION OPTIONS SURFACED

Accept · RFI · Reject at the asset level, with the matching cited STG clause and any past precedent.

OUTPUT THE SQP SIGNS

IRW Scope Assessment Record — SmartyGrants template, with the asset list, scope decision per item, and the SQP certification block.

IRW

Immediate Reconstruction Works — Finalisation

Finalise IRW delivery — works substantially complete within 12 months of the event, evidence audited line by line.

INPUTS THE ASSESSOR RECEIVES

- Council Finalisation package (cost data, completion photos, sign-offs)
- Labour ledger + GL transactions
- Completion photos and ITP / handover records per asset

WHAT RAAGENT AUTOMATES

- Allowable Time Limit check against the Cat B fact sheet
- Ledger reconciliation + on-cost analysis
- Completion-photo audit per scope item
- Per-DM damage and completion photo overlay

DECISION OPTIONS SURFACED

Per asset and per ledger line: **Eligible** · **Query** · **Ineligible**, with the cited STG / DAG / fact-sheet clause.

OUTPUT THE SQP SIGNS

IRW Finalisation Assessment Record — SmartyGrants template, with assessed amount per asset, evidence audit summary, and SQP certification.

EPAR

Essential Public Asset Restoration — Scope Approval

Approve the EPAR scope and Estimated Reconstruction Cost (ERC) against the EPAR Disaster Grant Guidelines (Dec 2025).

INPUTS THE ASSESSOR RECEIVES

- SmartyGrants EPAR Scope package (often 60+ DMs and 3,000 photos)
- Cost Estimation Tool workbooks per asset
- Pre-condition + damage evidence per DM
- Council-submitted ERC schedule

WHAT RAAGENT AUTOMATES

- CET workbook parsing against STG v2.1 unit rates
- Variance >5% flagging per line, plus contingency / PM / design / escalation band checks
- Photo bucketing per DM with EXIF + GPS quality flag
- Essential-public-asset eligibility cross-check vs the Eligibility Examples fact sheet

DECISION OPTIONS SURFACED

Per CET line: **Accept at CET** · **Accept applied** · **RFI for breakdown**. Per DM photo set: **Strong** · **Adequate** · **Insufficient**.

OUTPUT THE SQP SIGNS

EPAR Scope SQP Assessment Record — full GMS-DAF / 2025 v2.0 schema: Declarations · D1–D6 + C1–C4 · S1–S6. Auto-suffixed **_DRAFT** until every required tick is satisfied.

POST /api/submissions/<id>/epar-assessment-form

EPAR Essential Public Asset Restoration — Finalisation

Reconcile EPAR delivery against the approved ERC and pre-disaster function — recommend acquittal at the assessed amount.

INPUTS THE ASSESSOR RECEIVES

- Council Acquittal package with as-built cost and evidence
- Final-state completion photos per DM
- Variation and review history per the EPAR Variations fact sheet

WHAT RAAGENT AUTOMATES

- As-built vs ERC variance check per scope item
- Function-standards check against EPAR Function Standards & Alternate Solutions
- Re-damaged checks against the Re-Damaged fact sheet
- Final cost reconciliation per the Project Management & Other Costs fact sheet

DECISION OPTIONS SURFACED

Per scope item: **Acquit at ERC** · **Acquit at as-built** · **RFI on variance** · **Disallow**. SQP certifies overall recommendation.

OUTPUT THE SQP SIGNS

EPAR Finalisation SQP Assessment Record — same schema as Scope, with the acquittal amount, variance commentary, and final SQP certification.

Four lenses on the package. Each presents options. None decide.

PHOTO CHECKER

Three decision options per photo

- **Accept** · evidence sufficient
- **RFI** · scale / date / category needs clarification
- **Reject** · no damage shown or wrong asset

VISION MODEL

Claude Sonnet 4.5 audits the photo against the stated claim category + Photo Evidence Guidelines.

LEDGER CHECKER

Three verdicts per line + on-cost panel

Eligible · **Query** · **Ineligible** — each with the cited reason and evidence required.

ON-COST ANALYSIS

Applied rates vs RA-approved **33.08% FY24-25** and **34.62% FY25-26** per the Employee On-Costs Fact Sheet v4.

CET VERIFIER

Rate cross-check vs STG v2.1

Every CET line's applied rate compared to Cost Estimation Tool v2.1. Variance >5% flagged. Band checks: PM 6–12% · Design 4–10% · Contingency 10–20% · Escalation 4–12%.

Three per-line options: **Accept at CET** · **Accept applied** · **RFI for breakdown**.

DRFA EXPERT

Grounded chat — no hallucinations

Cites only the canonical EPAR reference set + DAG 2021 + DRFA 2018. For EW / IRW questions, explicitly notes when the canonical reference set has not yet been supplied.

Injects matching precedent into every answer (cited explicitly — never silent).

Authoritative for every EPAR Scope and EPAR Finalisation assessment.

TITLE	WHY IT'S CITED	SOURCE
NSW Essential Public Asset Restoration Guidelines (Dec 2025)	Primary EPAR reference. Cat B funding scope, essential-public-asset definition, eligibility tests, pre-condition certification, betterment cap, evidence rules, PM & on-costs, variation reviews.	nsw.gov.au/...
DRFA Standard Treatment Guide v2.1	Treatment / scope / unit-rate authority. The schedule the CET Verifier and STG matching engine compare against.	nsw.gov.au/...
Cost Estimation Tool v2.1 (external)	The published unit-rate workbook councils complete. RAAgent parses CET workbooks and cross-checks every line against this canonical set.	nsw.gov.au/...
EPAR Subcategories – Category B Restoration Works Fact Sheet	Defines the work categories the assessor selects against.	nsw.gov.au/166527
EPAR Eligibility Examples Fact Sheet	Worked examples the DRFA Expert quotes back when adjudicating borderline scope.	nsw.gov.au/166548
EPAR Evidence Requirements Fact Sheet	The Strong / Adequate / Insufficient call the Photo Checker surfaces is grounded in this fact sheet.	nsw.gov.au/166538
EPAR Function Standards & Alternate Solutions	Pre-disaster function definition; alternate-solution gates used at Finalisation.	nsw.gov.au/166528

The fact sheets that govern variation, redamage, on-costs & ERC.

TITLE	WHY IT'S CITED	SOURCE
EPAR Variations and Review Fact Sheet	Triggers and method for in-life scope variations. Used at Finalisation.	nsw.gov.au/166547
EPAR Re-Damaged Fact Sheet	Handling of re-damage events between Scope Approval and Finalisation.	nsw.gov.au/166543
EPAR Project Management and Other Costs Fact Sheet	Defines the PM / design / contingency / escalation bands the CET Verifier flags against.	nsw.gov.au/166673
EPAR Employee On-Costs (Opt-In Arrangements) Fact Sheet	The 33.08% (FY24-25) / 34.62% (FY25-26) RA-approved rates. Ledger Checker reconciles every applied rate against this fact sheet.	nsw.gov.au/171487
EPAR Establishing Estimated Reconstruction Costs (ERC) Fact Sheet	Methodology for the ERC the council submits – the basis of Scope Approval.	nsw.gov.au/167457
Allowable Time Limit – Category B Fact Sheet	The completion-window check the IRW & EPAR Finalisation flows enforce.	nsw.gov.au/166545
SmartyGrants Submission File Structure (Cat B / EPAR)	The canonical folder layout the ingest pipeline normalises against.	drfa-assess.com/...

The two cross-portal policy documents, and where EW & IRW reference sets stand.

TITLE	APPLIES TO	SOURCE
NSW Disaster Assistance Guidelines (DAG) 2021	All portals · eligibility framework, reasonable-cost provisions, definitions.	nsw.gov.au/18256
Disaster Recovery Funding Arrangements 2018 (Commonwealth)	All portals · the federal funding framework underneath the NSW DAG.	Internal mirror
Standard Treatment Tool v2.1 (external, IRW-only)	IRW Scope & Finalisation rate cross-check.	nsw.gov.au/...
EW reference set	EW Finalisation. PENDING – when RHE supplies the canonical EW reference pack, the Knowledge Base swaps in. Today the DRFA Expert relies on DAG 2021 + DRFA 2018 + STG v2.1 and explicitly recommends the assessor verify against the source.	Awaiting supply
EP12120000066 – Council Response 3	Worked precedent example bundled into the KB so the DRFA Expert can cite a real RFI cycle.	Internal mirror

KNOWLEDGE BASE DISCIPLINE

Every KB card carries an **applies-to** portal tag so EPAR references don't pollute EW questions. Every card has a one-click link to the authoritative source PDF. Every Expert answer cites by exact title.

What the DRFA Expert is allowed to say.

SYSTEM PROMPT RULES – NON-NEGOTIABLE

- ▶ EPAR answers ONLY cite from the EPAR reference set + DAG 2021 + DRFA 2018
- ▶ EW / IRW answers explicitly note that the canonical reference pack has not yet been supplied
- ▶ NEVER cite Form 306 – it is being phased out
- ▶ The OUTPUT is a completed EW Spreadsheet or EPAR SQP Record – not a free-form report
- ▶ Style: concise, precise, government-formal
- ▶ Present DECISION OPTIONS to the assessor – not a single recommendation

WHAT IT WILL REFUSE TO DO

- ▶ Invent guideline clauses · if unsure, recommend verifying against the source
- ▶ Mark an outcome on its own – every decision must come from an assessor click
- ▶ Cite a document outside the KB · the KB is the corpus
- ▶ Skip a citation · every claim names the fact sheet by its exact title

MODEL

Claude Sonnet 4.5 via the Emergent universal LLM key. 90-second timeout. Friendly 4xx surfaced on provider rejection (no opaque 500s reach the assessor).

The AI extracts and classifies. It does not write outcomes.

Where AI is invoked

- **Ingest classification** — heuristic + lightweight LLM tag on filename / header
- **PDF defect extraction** — single Claude Sonnet 4.5 call per package, on explicit assessor button
- **Photo Checker vision** — Claude 4.5 vision per photo, on explicit click
- **DRFA Expert chat** — Claude 4.5 streaming responses, grounded in KB + precedent corpus
- **Precedent embeddings** — used for semantic search of past assessor rationales

Where AI is NOT invoked

- **Outcome marking** — every Accept · RFI · Reject comes from an assessor click
- **Ledger arithmetic** — deterministic openpyxl streaming, not an LLM
- **Rate cross-check** — deterministic comparison vs CET v2.1 schedule
- **On-cost reconciliation** — pure arithmetic against the fact-sheet rates
- **Form generation** — direct openpyxl write of the SmartyGrants templates

WHY THIS MATTERS AT AUDIT

Every number on the EW Spreadsheet or EPAR SQP Record was either typed by a human, derived by deterministic arithmetic, or extracted by the AI and then explicitly accepted by an assessor click. There is no path by which an AI inference becomes a signed outcome without a documented human decision.

Institutional memory, made explicit.

Every assessor decision is recorded with the free-text rationale, indexed by job, by assessor, by date — and cited back inside the DRFA Expert's next answer.

What is recorded

- The line-item context (asset, defect, ledger line or CET line)
- The cited rule/clause
- The decision: **Accept** · **RFI** · **Reject** (or **Eligible** / **Query** / **Ineligible**)
- The free-text rationale typed by the assessor
- The assessor identity, job ID, and ISO timestamp

Mongo collection: **precedents** · API: GET /api/precedents

How it is used

- Vector search retrieves the closest past decisions for the current line
- The DRFA Expert receives them as context with explicit citations
- Citations are NEVER silent — every cited precedent is shown to the assessor with date, job, assessor name and rationale
- Browseable at /precedents — full corpus, filterable

WHY THIS IS NOT "THE AI LEARNING"

RAAgent does not silently fine-tune on council data. It surfaces the human decisions made on prior packages so the next assessor sees how a peer ruled — and decides their own line accordingly. Every precedent is auditable; nothing is opaque.

The DRFA Agent assists today. It becomes the SQP tomorrow.

Each phase shift is a governance milestone — Programme Director sign-off, NSWRA notification, public statement update on the platform. Never a silent upgrade.

PHASE 1 · TODAY

The Agent assists. The SQP decides every line.

The platform classifies, parses, reconciles, surfaces options. The assessor clicks. The click is what is written to the artefact.

PHASE 2 · NEXT

Defaults on the easy ones. Assessor confirms in one click.

On high-confidence, low-risk lines the Agent pre-selects the recommended outcome. The SQP one-click confirms or overrides. Programme Director sign-off and NSWRA notification before each portal turns on.

PHASE 3 · ENDGAME

The Agent IS the SQP where it has out-performed the human baseline.

Category by category, against a published audit benchmark, the Agent assumes the assessor mark. Humans handle escalations, governance and audit.

THE THRESHOLD TO ADVANCE A CATEGORY

RHE publishes an annual benchmark of agent-vs-assessor agreement, RFI accuracy, and audit-survivability. A portal category advances only when the agent has beaten the human baseline for two consecutive quarters AND NSWRA has sighted and signed the benchmark.

Built to be defensible in an audit.

Sign-off chain

- Every assessment artefact (EW Spreadsheet, EPAR SQP Record) carries an SQP name + role + date
- The EPAR Record is auto-suffixed _DRAFT until every required tick (D1–D6, C1–C4 \geq 10; S1–S6 \geq 6) is satisfied
- QS / Engineer flag on every EPAR Record
- Per-portal SLA timer + Resume Jobs panel – every job traceable to a named assessor

RHE Assessor framework alignment

- RHE Project Consulting operates under the NSW DRFA assessor framework
- RAAgent is an internal RHE tool – not an NSWRA-operated system
- All outputs are RHE-signed SQP records that NSWRA receives via SmartyGrants
- Phase transitions require Programme Director sign-off (RHE) AND NSWRA notification

Domain-locked, audit-traced, no shared state.

Access

- Domain-locked sign-in to @rheprojectconsulting.com.au + explicit assessor allow-list
- Standalone tools (Ledger Checker, Photo Checker, CET Verifier) available without sign-in – public, no data persistence on first use
- NSW Gov SSO on the roadmap (replaces the domain allowlist when ready)
- No multi-tenancy bleed: a council's package never appears in another council's precedent search unless explicitly de-identified

Secrets & transport

- Zero hardcoded secrets · all credentials in env files only
- HTTPS-only · TLS terminated at the Emergent ingress
- CORS managed via env · production locks down origins
- LLM calls wrapped in a 90 s timeout with graceful 4xx on provider rejection

WHERE TO READ THE LIVE SECURITY & AI STATEMENT

<https://drfa-assess.com/security> · also linked from the header on every page

Every click traceable. Every artefact reproducible.

What is logged

- Every ingest job: source URL, timestamp, classifier confidence, milestone log
- Every assessor decision: line context, cited clause, decision option, free-text rationale, ISO timestamp, assessor identity
- Every form generation: assessor name + role, date certified, draft / final flag
- Every DRFA Expert answer: cited fact sheets + cited precedents, both surfaced to the assessor

Data residency

- Application runs on Emergent's ARM cloud image
- MongoDB managed by Emergent
- Ingested packages stored under /app/backend/work/<job_id> on the application volume; cleared per disk-pruning policy
- Verified-tree snapshots persisted to Mongo so the assessor view survives pruning
- No third-party analytics or telemetry beyond the platform's own logs

Two-environment discipline. Test before prod.

Environments

- ▶ **Preview (dev):** internal RHE testing — every code change lands here first
- ▶ **Production:** `drfa-assess.com` — promoted only by explicit RHE deploy
- ▶ Two replicas, 250m CPU, 1 GiB memory each, behind Cloudflare

Backups & recovery

- ▶ MongoDB managed-service snapshots
- ▶ Verified-tree snapshots in Mongo (survive disk pruning)
- ▶ Generated artefacts (videos, packages) shipped via the frontend build image (versioned per deploy)

Testing posture

- ▶ 43+ backend pytests at `/app/backend/tests/`
- ▶ Playwright frontend smoke tests across all five portals
- ▶ Deployment-readiness scan (env, supervisors, CORS, query optimisation) before every promote
- ▶ Static-analysis linters block any commit with syntax / undefined / unused-import errors

Observability

- ▶ Per-job ingest milestones polled by the UI in real time
- ▶ Backend access logs aggregated `/var/log/supervisor/` under
- ▶ Live SLA timer panel on every assessor dashboard

The SQP signs. RAAgent did the data work. NSWRA receives an auditable record.

RAAgent ingests a SmartyGrants disaster-recovery package, classifies every file, reconciles the labour ledger against the NSWRA on-cost fact sheet, buckets damage and completion photos per Damage Memo, presents the assessor with structured decision options at every line — with the matching past precedent already cited — and generates the canonical NSWRA artefact at sign-off.

The AI never marks an outcome. The assessor clicks. The click is what is written.

THE PLATFORM SHIPPING TODAY

Five portals · sixteen authoritative references · four assistant tools · precedent corpus · audit-grade outputs in SmartyGrants format · live at drfa-assess.com.

LIVE NOW

Where to go from here.

SURFACE	URL
Showcase landing	https://drfa-assess.com/showcase
Production portal	https://drfa-assess.com
EW Finalisation walkthrough video	drfa-assess.com/media/EW_Finalisation_Walkthrough.mp4
Lead marketing video	drfa-assess.com/media/RAAgent_Lead_Marketing_Video.mp4
Security & AI statement	drfa-assess.com/security
Precedent corpus	drfa-assess.com/precedents
Knowledge Base	drfa-assess.com/knowledge-base
This deck (live, print-friendly)	drfa-assess.com/api/presentation

RHE Project Consulting · drfa-assess.com · Feb 2026

