



ToM Assurance Suite

AI-native assurance intelligence for requirements, evidence, interfaces, and governed review of infrastructure projects.

Newport Resonance | Information Deck | 2026



Context Retention



Standards Intelligence



Zero-Copy Custody



Auditable



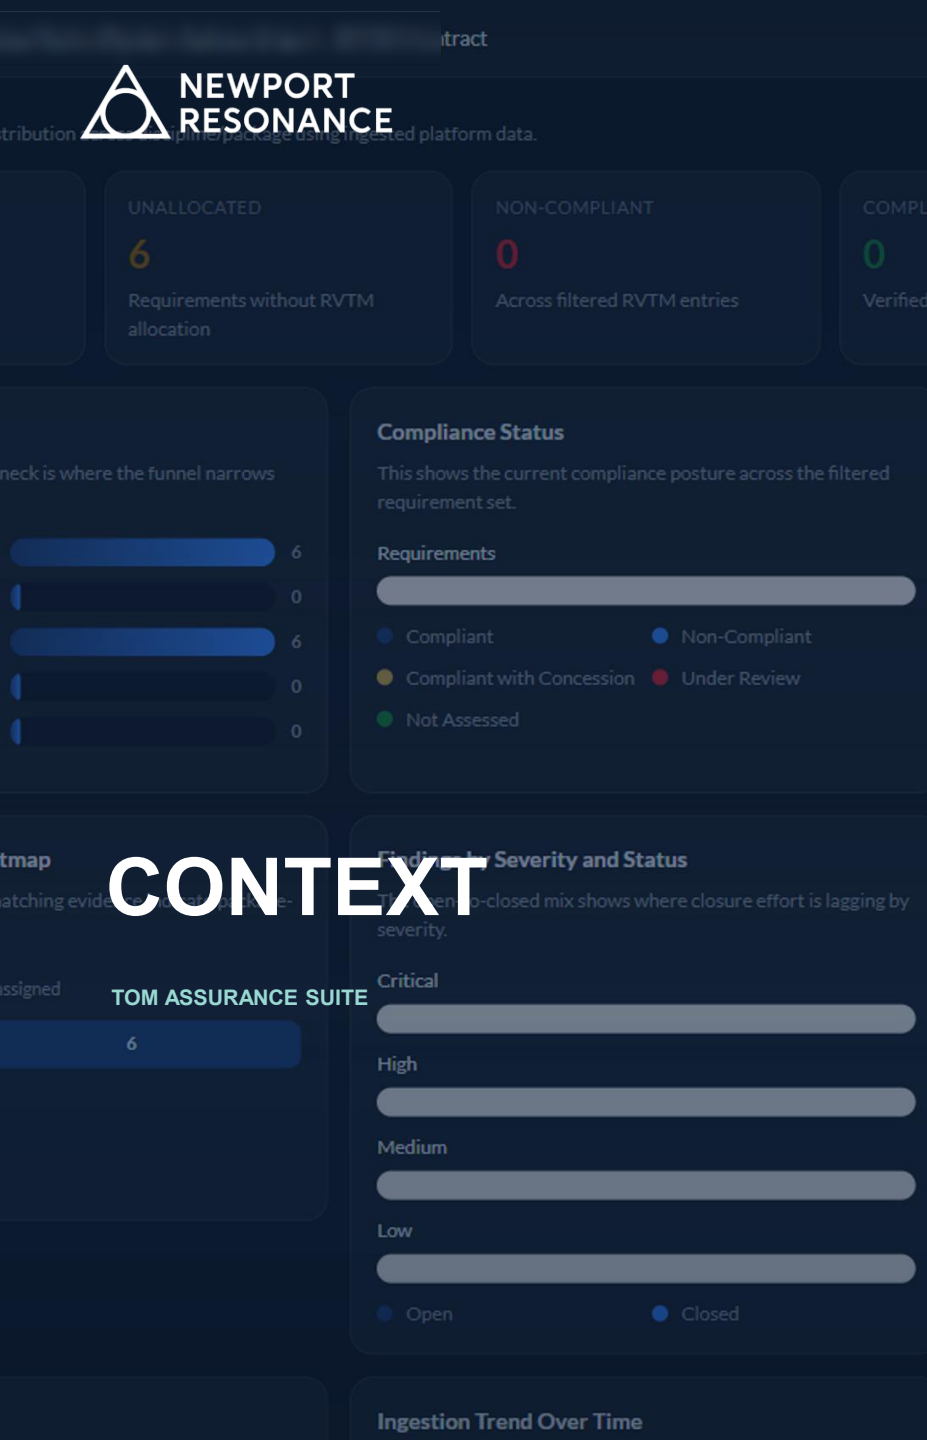
Real-time

INDEX

TOM ASSURANCE SUITE

Contents

01	Executive context	Pages 3-4	Why assurance breaks and what users need to prove.
02	Company and product position	Pages 5-7	About Newport Resonance, ToM position, and the assurance operating loop.
03	Differentiated architecture	Pages 8-10	ToM's innovative design, governed continuity, source custody, review state, and data boundary.
04	Core capabilities	Pages 11-19	Source intake, requirements, evidence, RVTM, interfaces, findings, reports, assistant, initiatives.
05	Proof, and value	Pages 20-21	Where ToM fits, current proof points, and user outcomes.
06	Pilot and next step	Pages 22-23	Pilot structure, engagement paths, and contact details.



03 / CONTEXT

The assurance problem is no longer manual. It is automatic and progressive.

ToM Assurance Suite keeps project requirements, design evidence, decisions, and reports connected as the project develops.

What ToM does

ToM Assurance automatically reads the contract requirements that govern a project:

- specifications,
- Australian Standards,
- third-party agreements,
- planning obligations,
- interface agreements

It works out which obligations are relevant to the actual scope and checks them against the evidence that has been uploaded or received through approved document-control pathways, including platforms such as **Aconex** or **TeamBinder** where integration is approved. The ToM assurance Suite has advanced intelligence capabilities to read, understand and infer evidence from:


- design reports,
- Engineering drawings,
- test records.


Live compliance matrix

The product then maintains a live requirements compliance matrix, or RVTM, where every requirement carries source references, evidence links, confidence, review state, project-stage context, and risk signals. A concept-stage project lifecycle is treated differently from a missing construction or AFC gate requirement, and an internal package issue is separated from a third-party or “by-others” interface risk.

Why it matters

Current infrastructure projects do much of this work manually across large design/assurance teams, spreadsheets, document registers, and late-stage review cycles. That consumes months of effort and still leaves gaps when evidence changes progressively. Many large projects spend \$ millions on independent certifiers and ToM assesses compliance automatically in real time. ToM keeps the source chain alive so a dashboard number, report conclusion, compliance row, or finding can be audited back to the clause, drawing, test record, or reviewer decision behind it.





CONTEXT

04 / CONTEXT

Assurance breaks in the gaps between systems, not inside one file.

Many projects change faster than manual traceability can keep up. Multi-discipline coordination, design interface and construction driven changes that miss a downstream impact.

Manual assurance does not scale

Large rail and infrastructure projects carry tens of thousands of requirements and evidence fragments across registers, drawings, models, reports, meetings, and workflow systems. Teams spend assurance effort reconstructing links instead of resolving risk.


Risk appears in the handoffs

When drawings supersede assumptions, AFC gates move forward, interfaces change, or evidence arrives unlinked, dashboards can show progress while hiding stale conclusions, missing source basis, and unresolved review decisions.

The change ToM makes

ToM turns assurance from manual collation into a governed evidence workflow. It keeps requirements, evidence, confidence, maturity, interface risk, reviewer decisions, and report basis connected so teams focus on exceptions that affect cost, programme, certification, accreditation, and handover.

WHERE ASSURANCE BREAKS



The risk is in lost relationships, not missing files.

NEWPORT RESONANCE

TOM ASSURANCE SUITE

05 / NEWPORT

Newport Resonance builds governed intelligence for high-consequence projects.

We combine engineering discipline with proprietary ToM intelligence infrastructure to make complex project decisions source-backed, auditable, and reviewable.

Engineering-led intelligence

Newport Resonance was founded to apply engineering discipline to AI-enabled decision support. In infrastructure assurance, every important conclusion needs a source basis, authority context, confidence basis, and visible review path.

The ToM intelligence system

ToM provides the governed intelligence layer behind the ToM Assurance Suite and other products. It combines long-term memory, policy controls, source traceability, confidence scoring, bounded adaptation, and reviewer state so outputs can be inspected before they are relied on.

Responsible development


Our products are designed for environments where unsupported recommendations create real risk: *safety, programme, commercial, accreditation, and handover risk*. The goal is not faster answers alone; it is faster decisions with visible evidence, limits, and accountability.

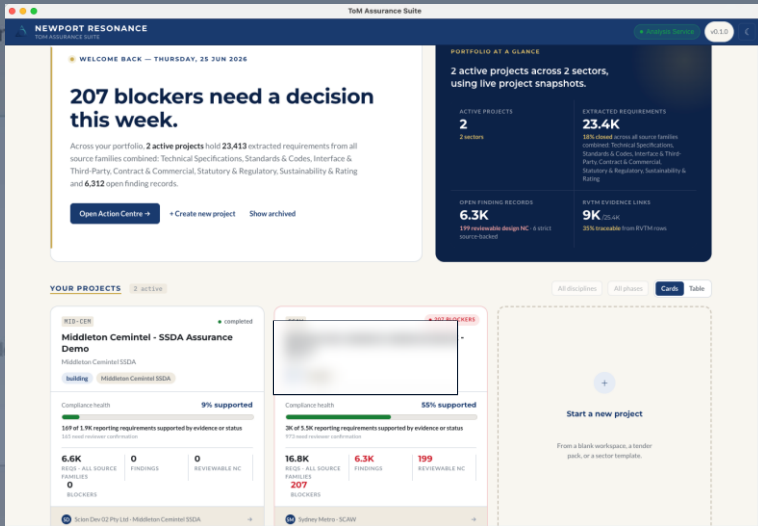
The Newport Resonance position is deliberately different from broad AI automation and typical “LLM wrapper” products that apply a RAG memory with vector database, but still derive the bulk of its intelligence through a general-purpose LLM.

The company is building intelligence infrastructure for work where a recommendation is not valuable unless the receiver can inspect the basis for it. That means including confidence, limitation states, and human review as part of the product, ultimately meaning that **the intelligence system learns and adapts from a feedback loop as it is used**.

Capability detail

- **Source-backed reasoning:** conclusions always remain linked to evidence.
- **Policy-guarded outputs:** reliance limits and review gates are explicit.
- **Confidence visibility:** uncertainty and basis are shown, not hidden.
- **Bounded adaptation:** the system improves from governed feedback.
- **Human review:** important decisions remain inspectable and accepted.





PRODUCT

TOM ASSURANCE SUITE

No sustainability, environmental ratings, or ESG obligations loaded yet.

[Open Obligation Sources](#)



Not loaded

CONTRACT & COMMERCIAL

No deed, contract, or commercial sources loaded yet.

[Open Obligation Sources](#)

Recent Activity

0 / 0 (0%)

06 / PRODUCT

An assurance intelligence overlay for requirements, evidence, traceability, and compliance.

Designed to sit above existing systems, not replace them on day one.

Overlay, not replacement


ToM is designed to sit above the systems a project already uses: requirements tools, document-control platforms, design repositories, spreadsheets, BIM/CAD sources, project controls, and reporting templates. The first value case is not a disruptive migration. It is recovering the assurance meaning across the systems that already hold the proof.

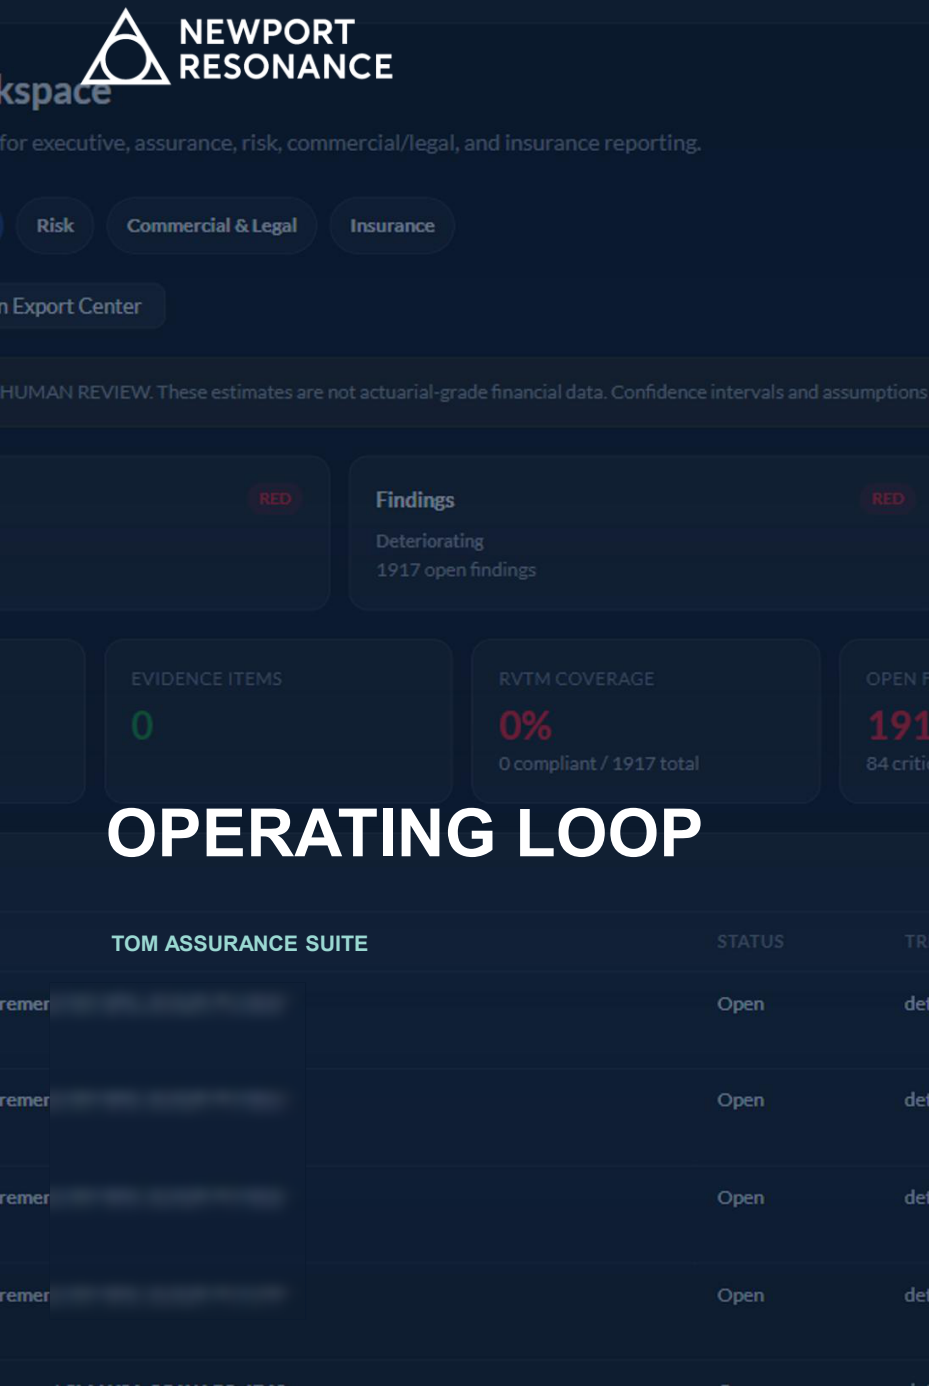
The core product output

The operating output is a live, auditable compliance matrix and assurance record. ToM shows what is required, which sources govern it, what design or engineering evidence supports it, what confidence the system has, what remains unresolved, what needs human review, and where the risk sits in the project lifecycle.

Adoption path

A customer can begin with one package, standards problem, RVTM audit, interface set, source-readiness review, or pre-issue/AFC gate. Once the evidence workflow is proven, the same model can expand by discipline, package, project, sector, or enterprise governance model.





07 / OPERATING LOOP

The same connected record powers dashboards, RVTMs, reports, interface packs, and assistant answers.

ToM turns project source material into a live, reviewable assurance record.

Read, classify, connect

The operating loop starts with source intake. ToM classifies documents by role and context, extracts requirements and source facts, identifies relevant clauses and obligations, and creates structured assurance objects that can be reviewed rather than treated as one-off AI answers.

Check evidence against obligations

The system matches design reports, drawings, calculations, test records, interface documents, notices, and other evidence to the requirements they support. It separates strong source-backed evidence from weak signals, stale references, missing information, and rows that need reviewer judgement.

Keep updating as the project changes

As new evidence is uploaded or transmitted through approved platforms, the assurance record updates. Dashboards, RVTM rows, findings, interface artefacts, report packs, and assistant answers use the same connected source chain, so change does not force the team to rebuild traceability from scratch.



ARCHITECTURE

TOM ASSURANCE SUITE

Not a chatbot on a contract

A chatbot can retrieve passages, summarise a document, and draft a fluent response. That is useful, but it does not by itself create an assurance regime. The hard question is whether a claim is current, relevant to the works, supported by evidence, phase-appropriate, accepted for reliance, and auditable under review. The ToM assurance suite infers information from engineering drawings, not just read information verbatim.

Governed AI architecture

ToM wraps the language model in a machine-learning governance layer.

Outputs and responses are constrained to project records, source roles, evidence links, confidence scores, review states, and policy gates.

If ToM cannot substantiate a requirement, it should flag a gap or review item rather than inventing an answer.


User consequence

This structurally reduces the hallucination risk that makes engineering teams wary of generative AI. The value is not that the system sounds confident. The value is that every compliance decision can be traced back to a clause, drawing, test record, evidence file, reviewer decision, or clearly marked limitation.

08 / ARCHITECTURE

An innovative System Architecture to minimise AI hallucination.

The difference is not whether an LLM is used. The difference is what governs the source, evidence, confidence, and review trail. ToM is not just an LLM Wrapper, it uses the LLM as part of a reasoning loop.



ARCHITECTURE

TOM ASSURANCE SUITE

09 / ARCHITECTURE

ToM turns the model into a governed assurance system.

The model reasons, but ToM preserves project memory, evidence state, policy gates, confidence basis, and reviewer control around it.

Project continuity beyond the prompt

ToM uses an AI LLM as one reasoning component inside a governed assurance product. The surrounding architecture holds the project memory, source boundaries, requirement records, evidence links, reviewer states, confidence basis, policy gates, and learning controls that a normal prompt or retrieval workflow does not maintain.

Evidence change becomes assurance impact

If a specification, drawing, design report, test record, or notice changes, ToM can interpret the change against the assurance record. It can identify affected requirements, stale evidence links, findings needing review, interface artefacts that should be updated, and report conclusions that should no longer be relied on.

Improvement under governance

The system is designed to improve from reviewed project patterns. The more it is used, the better it can recognise recurring compliance and risk signals, but learning remains governed through admission, provenance, quarantine, supersession, decay, and human review rather than uncontrolled accumulation of raw project data.



Major infrastructure projects spend tens of millions confirming contract compliance. 30,000+ requirements per project, each confirmed manually.

Trust starts with **intelligence**, not data management.

☒ **Thinks like your engineers**
Specialist domain review against standards

☒ **Understands obligations**
Parses obligations

☒ **Learns continuously**
Cross-project intelligence, not keyword matching

TRUST

TOM ASSURANCE SUITE
Validation Test Suite

VAL-001 Rail
Created Mar 31, 2026

Created Mar 13, 2026

+
Create New Project

10 / TRUST

Trust starts with zero-copy custody, not after-the-fact controls.

ToM keeps project evidence anchored to governed source systems while applying policy gates, reviewer state, and confidence limits around every assurance output.

Zero-copy source custody


ToM references governed project evidence without turning copied files, prompts, or chat history into the system of record. Source boundaries, reliance limits, and auditability remain explicit.

No free-form assurance opinions

Findings must be anchored to a source: *clause, drawing, calculation, test record, notice, interface document, evidence file, or reviewer decision*. Unsupported conclusions become gaps, caveats, or review items.

Visible reliance state

Users can distinguish draft synthesis, review-only signals, evidence-backed findings, accepted decisions, and unresolved gaps. Confidence means showing what is known, what is uncertain, and who accepted the result.



STANDARDS & INFERRED COMPLIANCE

11 / STANDARDS INTELLIGENCE

Persist standards once, infer what applies, then test compliance with evidence.

ToM keeps the full standards library visible while inferring project-specific applicability and testing active clauses against drawings, records, and engineering evidence.

Reusable standards library


Key standards, codes, approvals, and third-party obligations are ingested once, structured into persistent clause tables, and reused across projects. Teams can select known sources rather than rebuilding standards checklists from scratch.

Applicability is inferred, not only user selected

ToM compares standards clauses against project scope, assets, materials, packages, design stage, geography, discipline, and interface context. It proposes active, uncertain, review-required, or inferred-not-relevant clauses with confidence and source-backed reasons.

Compliance reasoning moves beyond evidence matching

The product is being developed to infer standards compliance from engineering evidence, not just uploaded declarations. For example, where reinforcement is shown on a concrete beam drawing, ToM can compare likely reinforcement area against AS 3600 requirements and flag whether the clause appears satisfied, uncertain, or in need of review.



*Example standards can be preloaded to the ToM Assurance Suite



Type to filter

12 / 23

EVIDENCE

TOM ASSURANCE SUITE

13 / EVIDENCE

Evidence proves sufficiency. RVTMs remain comparators until reviewed.

ToM separates source-backed evidence, live trace state, reviewer decisions, and imported spreadsheet claims so assurance progress does not hide stale or unsupported conclusions.

Evidence is not just a file

ToM assesses whether evidence is relevant, current, phase-appropriate, strong enough, and linked to the right requirement with a clear method and confidence basis.

The RVTM becomes live assurance state


Rows can show governing requirement, linked evidence, confidence, reviewer decision, stale references, interface risk, and closure readiness as new evidence arrives.

Imported RVTMs are audit evidence, not truth

Client or legacy RVTMs are useful for comparison, migration, stale-record detection, and source acquisition planning. They do not automatically become accepted assurance records.

Pre-issue and AFC checks

This supports design-gate assurance before deliverables are issued or marked Approved for Construction. Requirements without adequate evidence are surfaced at the gate, where they can still be resolved through design correction, RFI, review, or source acquisition, rather than discovered on site as rework, claims, or concessions.



BIM may show a clash; ToM captures the interface obligation, ownership, evidence basis, maturity, and agreement path.

Internal Package Boundary External Contract External / Third Party System

Now Not Due Yet Required Missing Uncoordinated Source Conf

INTERFACES

TOM ASSURANCE SUITE
Documents & Evidence

turity risk separates due-now missing information from future design-stage evidence.

LEAD INTERFACE MATCH INTERFACE INTERFAC

Type to filter Type to filter Type to

Durability and Whole of Permanent Road Package Finding t

ToM finds and governs interfaces that ordinary coordination misses.

Interfaces are inferred from requirements, drawings, reports, and design evidence, then converted into source-backed coordination records to mitigate project risks.

Inferred, not only declared


ToM reviews requirements, design drawings, technical reports, evidence records, and project correspondence to identify interfaces created through the design process, not just interfaces manually entered by users.

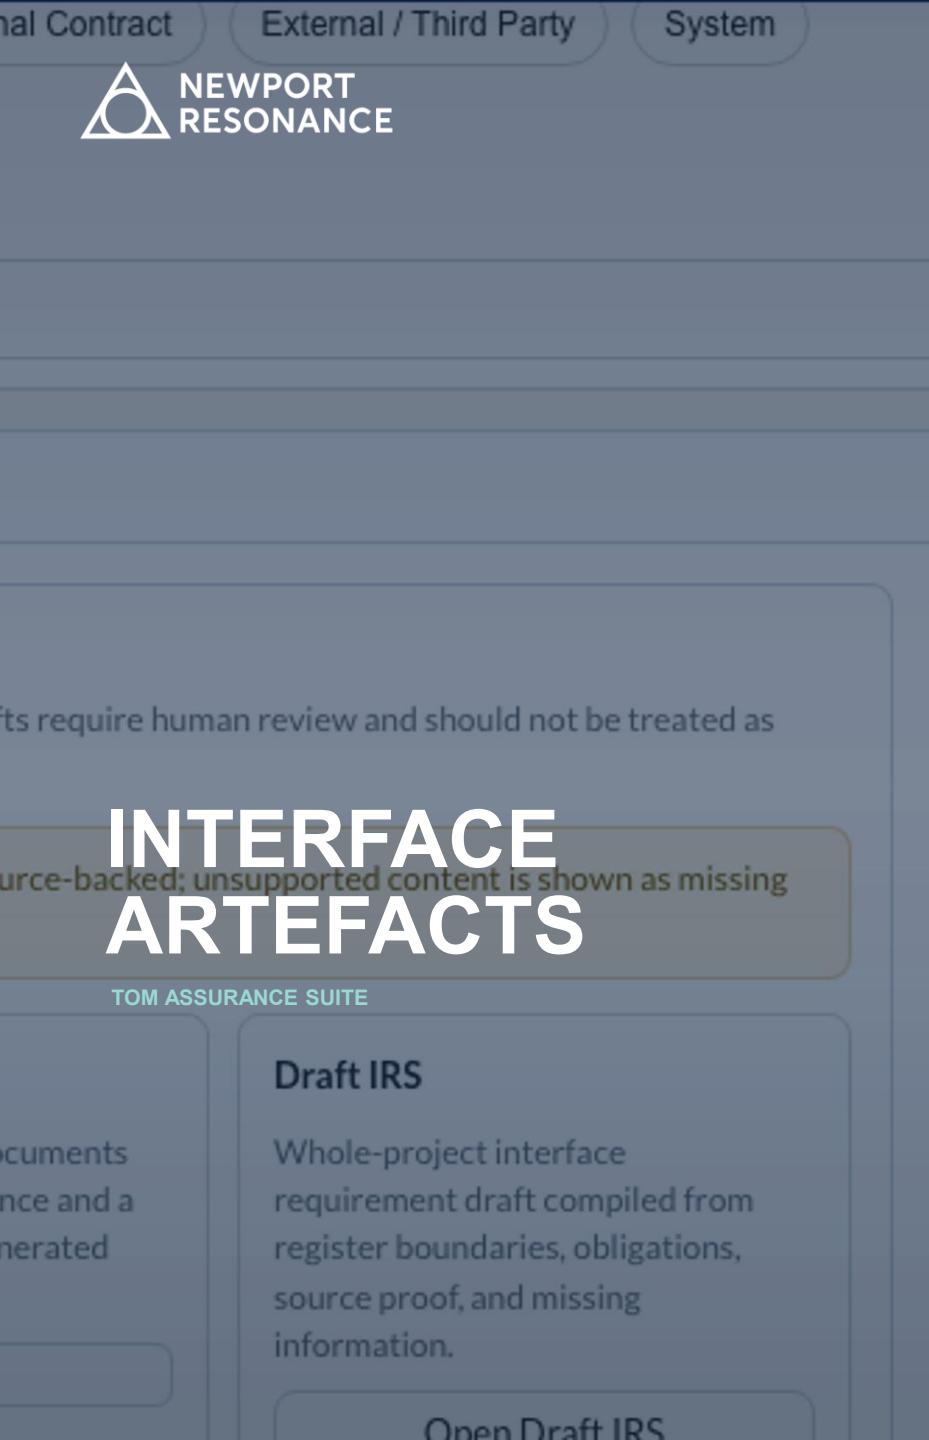
Covers real project boundaries

Interfaces are classified across intra-package, inter-package, third-party, system, staging, handover, and external contract boundaries, with source basis, confidence, maturity, and risk state.

From discovery to agreement

Each interface can be captured in the register with lead party, match party, obligations, evidence, due-now status, and unresolved decisions. ToM can then generate Interface Requirement Specifications (IRS) and Interface Control Documents (ICD) drafts so coordination moves from hidden risk to documented agreement.





Turn inferred interfaces into reviewable coordination artefacts.

ToM generates draft IRS and ICD packs from source-backed interface records, while preserving missing information, caveats, and reviewer control.

Interface records become artefacts

ToM converts inferred and registered interface relationships into structured review packs that help lead and match parties understand scope, responsibility, source basis, maturity, risk, and next action. A BIM model does not capture details like structural load transfer capacity at an interface.

Missing information stays visible

Draft artefacts distinguish known facts from unsupported fields. Missing demarcation, drawings, acceptance criteria, agreement state, or party responsibilities are carried into a missing-information register instead of being silently filled.


Review before reliance

Draft IRS and ICD outputs are not issued contractual documents. They accelerate coordination by giving reviewers a source-backed starting point, with caveats, provenance, and QA checks preserved.

Capability detail

- Draft IRS/ICD generation from interface register facts
- Source registers and missing-information registers
- Boundary figures and coordination summaries
- QA caveats, provenance, and reviewer acceptance





Find reviewable design non-compliance signals before they become disputes.

ToM turns notices, source facts, design evidence, and weak assurance signals into governed findings without promoting candidates to final determinations.

Signals stay source-backed


ToM identifies candidate findings from technical specifications, drawings, reports, notices, interface records, and evidence gaps. Each signal retains its source chain, confidence basis, affected requirement, and project context.


Candidates are not determinations

A candidate non-compliance, notice-linked issue, evidence gap, or review blocker remains reviewable until accepted. The system separates weak signals from confirmed findings so assurance teams can act early without overstating the record.

Action routing preserves judgement

Findings move into review queues, escalation paths, report packs, or closure workflows with reviewer state preserved. The product helps teams focus on issues that may affect cost, programme, certification, accreditation, or handover.





17 / DASHBOARDS

One operating picture from project evidence to executive decision.

Dashboards, report packs, and exposure indicators inherit the same source-backed assurance record, so progress, blockers, caveats, and commercial risk stay connected.

Live assurance state

ToM dashboards show what is known, weak, blocked, processing, review-ready, or unresolved. Metrics remain entry points into the source record, not detached presentation numbers. But the platform doesn't just surface a risk – it assigns owners to issues so that the feedback loop is shortened. The intent of ToM Assurance is to ensure risks are known by the right people at the right time, not hidden in an appendix of a report.

Reports preserve the evidence basis

Report packs inherit source registers, missing-information registers, confidence basis, limitations, and reviewer state so executive narratives do not outrun the evidence.

Exposure remains caveated

Exposure indicators help leaders see where assurance gaps may create cost, programme, certification, accreditation, or handover risk. They are draft, confidence-banded, and review-dependent.



candidates:

$$0.51g N^*/(Ag f c \text{ [?]}) = 0.0 \text{ } 0.31g \text{ Walls } N^*/(Ag f c \text{ [?]})$$

E: For intermediate values of $N^*/(A f \text{ [?]})$, linear

etermined from the design creep coefficient of

c In the absence of more accurate calculations,

as- $s \frac{1}{(1 \text{ [?] } N^*/ \text{ [?] } N)} \dots 10.4.3(1) c$ where the

ated for c each column in accordance with

ToM ASSISTANT

TOM ASSURANCE SUITE

18 / ASSISTANT

Ask project-grounded questions without losing provenance.

ToM Assistant answers against connected project sources with references, confidence, caveats, and next actions – not a made up or hallucinated response.

Governed entry point

ToM Assistant gives engineers a faster route into the project record, but it is not a free-floating chatbot. It answers against ingested project material, requirements, evidence links, review state, findings, standards posture, interface records, and report context.

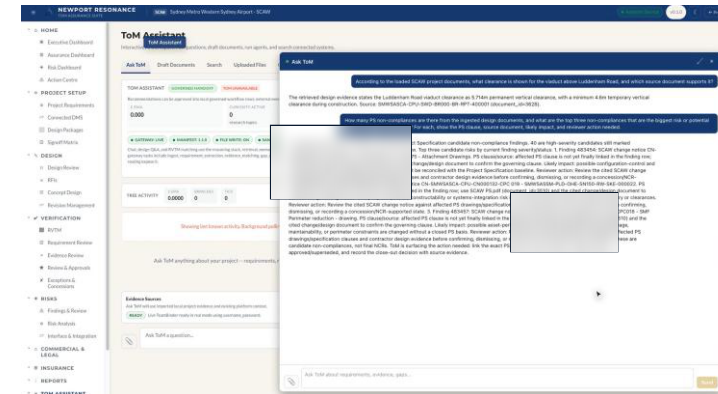
Need to find a technical answer from thousands of design drawings or reports? – the assistant will tell you with document references.

Every answer carries its basis

Useful answers show source references, confidence, limitations, and next action. A user should know whether the response is navigation help, draft synthesis, evidence-backed advice, or a review-required issue. Unsupported questions should produce caveats or gaps, not confident guesses.

Practical use

The assistant can help users ask which requirements lack evidence, why an RVTM row is uncertain, which clauses may apply, what interface obligations are unresolved, or what sources support a report statement. It remains a gateway into the assurance record rather than a parallel decision-maker.



GOVERNED HANDOFF
[Change m](#)

ance evidence and governed tom_master intelligence results.

yet available.

INITIATIVE

TOM ASSURANCE SUITE

19 / INITIATIVE

Let the platform surface the next assurance move.

Initiatives are reviewable recommendations, not uncontrolled autonomous execution. External signals stay read-only and evidence-gated.

From dashboard to action


ToM Initiative turns the assurance record into a prioritised work queue. It can surface blockers, weak evidence, source-quality issues, stale RVTM rows, quality trends, interface pressure, accreditation gaps, commercial exposure signals, and recurring compliance patterns.

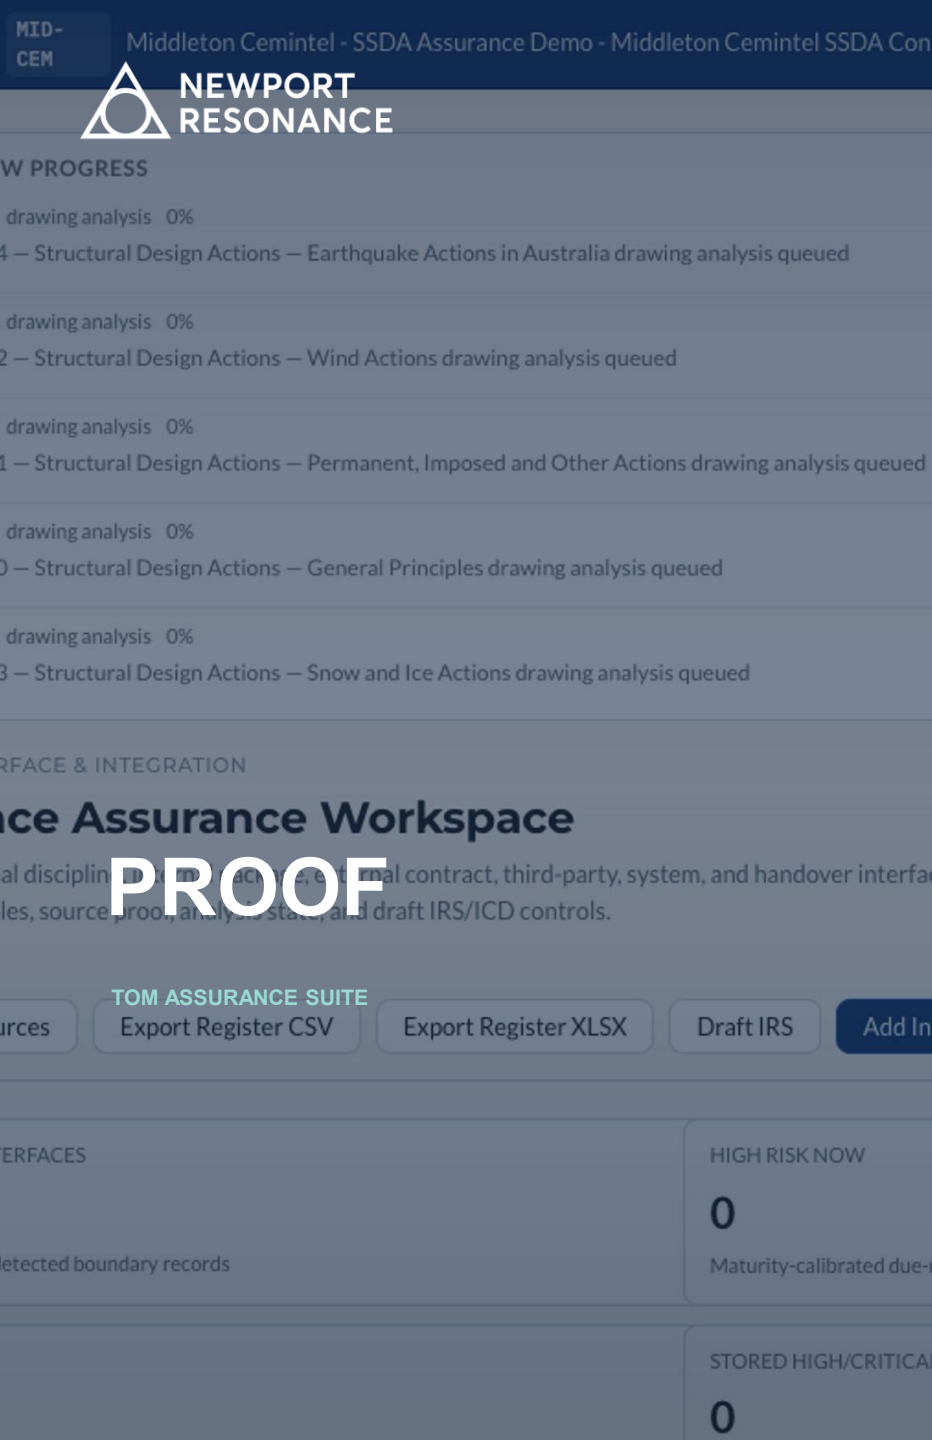
External signals stay governed

Where approved, ToM can check read-only external context such as standards updates, public authority material, operator requirements, planning conditions, network-owner obligations, or permitted document-system observations. It can browse the web for information on potential macro-economic impacts to projects. External information remains evidence-gated, policy-routed, and reviewed before reliance.

Self-learning with review

The distinctive value is the feedback loop. Reviewer decisions and high-quality project patterns can improve future recommendations, but learning candidates are admitted, quarantined, superseded, decayed, and reviewed under policy. The system gets better at spotting the project risks it has been taught to recognise without becoming uncontrolled automation.





20 / PROOF

Validated on real project assurance records.

ToM Assurance Suite has been developed and tested against representative building and infrastructure source sets with known requirements, certified design documentation, and existing assurance records.

Real source sets, not synthetic examples

Development has used project-scale contracts, specifications, drawings, design reports, evidence records, interface registers, and RVTM-style assurance data from representative building and transport infrastructure projects.


Benchmarked against known records

The system was tested against known requirements, final certified design documentation, and project assurance measurements to compare whether ToM could identify obligations, evidence links, compliance signals, and interface risks with greater detail and speed.

More than document retrieval

Testing focused on whether ToM could reason across source material: extracting requirements, assessing evidence sufficiency, identifying interface obligations, detecting stale or missing assurance links, and producing reviewable outputs with source basis and caveats.





Earlier clarity, stronger review quality, and less assurance drag.

The value is not weaker governance. It is better evidence before decisions are made.

Change the economics of assurance

The aim is materially lower assurance overhead without weakening governance. ToM automates much of the collation, evidence matching, RVTM maintenance, source checking, and gap detection that currently consumes large assurance teams and months of effort. This makes assurance activities more feasible across broader scope for better project risk management.


Earlier issue detection

Non-compliance, weak evidence, stale RVTM rows, missing third-party obligations, unresolved interfaces, and AFC gate issues can be surfaced while there is still time to act. Earlier visibility reduces late-stage contract surprises, concession pressure, downstream rework, RFIs, defects, and claims.

Key user outcomes

- Lower assurance cost and headcount by shifting specialists from collation to judgement.
- Pre-issue and AFC checks before non-compliance reaches site.
- Tighter commercial control through obligation, cost, variation, and claim-risk linkage.
- Interface and third-party requirement visibility across packages, operators, utilities, authorities, and adjacent contractors.





The next step is a real project walkthrough.

Bring a bounded source pack and the assurance decision that matters.

Run the demo on real sources

The best next step is a practical session with the assurance, systems-integration, design, or commercial lead who owns the problem. Bring a real source pack and one decision that matters: a requirements subset, RVTM audit, evidence gate, standards applicability problem, interface register, or report pack.

Prove the workflow

The walkthrough should show how ToM reads the governing requirements, identifies which obligations are relevant, checks available evidence, scores confidence, flags gaps, and produces a live reviewable matrix or pack with source references and reviewer actions.

Choose the value case

A useful first pilot should be narrow enough to run quickly and important enough to justify rollout. The fastest value usually comes from a known assurance pressure: stale RVTM rows, pre-issue/AFC compliance checks, interface risk, third-party obligation tracing, or executive reporting that currently requires manual reconstruction.

CONTACT

TOM ASSURANCE SUITE

23 / CONTACT

Questions or interested in a pilot?

We can walk through a demo, confirm product fit, and shape the right evaluation path. We are able to customise the interface for your business needs.

EMAIL info@newportresonance.com

WEB www.newportresonance.com

CONTACT **Ken Morkaya**

PHONE 