RESEARCH PARTNERSHIP DISCUSSION PACK

Universal Readiness Engine Platform Specification

Configurable evidence orchestration for domain-specific, human-reviewed readiness workflows

One engine. Domain-specific readiness applications. Evidence first. Human judgement retained.

PREPARED BY	Axiom Systems Pty Ltd
PLATFORM	**Universal Readiness Engine**
ISSUE DATE	19 August 2026
STATUS	NON-CONFIDENTIAL OVERVIEW

Purpose. This document is a non-confidential discussion aid for prospective research, industry and commercialisation partners. It does not constitute product certification, regulatory opinion, funding commitment or offer of legal terms.

© 2026 Axiom Systems Pty Ltd. All rights reserved.

Document map

Note: Open in Microsoft Word and use References → Table of Contents if an automatically paginated table of contents is required. This discussion version uses a concise document map to remain portable.

Section	Purpose
1	Executive specification
2	Problem model and design principles
3	Core operating model
4	Logical architecture
5	Canonical information model
6	Workflow and state model
7	Evidence intake, extraction and provenance
8	Human review and bounded AI assistance
9	Configuration and verticalisation
10	Integration model
11	Security, privacy and audit boundary
12	Non-functional requirements
13	Pilot deployment model
14	Demonstrated maturity and claim boundary
15	Candidate domain implementations
16	Research and validation agenda
17	Limitations and non-claims
18	Glossary and evidence basis

© 2026 Axiom Systems Pty Ltd. All rights reserved.

1. Executive specification

The Universal Readiness Engine (URE) is a configurable orchestration and assurance layer for work in which a team must determine whether a project, case, asset, organisation or intervention is ready to progress. It combines a defined target state, contextual evidence, configurable requirements, visible gaps, human review, owned actions and a traceable outcome record.

> **Core proposition**
>
> URE does not replace professional judgement or the systems that hold authoritative data. It gives teams a structured way to connect evidence to a readiness decision, preserve provenance, expose uncertainty and drive unresolved matters to closure.

Dimension	Description
Primary user outcome	A shared, reviewable answer to: What is known? What is missing or uncertain? Who must decide? What must happen next? Is the agreed target state ready to progress?
Primary system role	Orchestration across briefs, documents, data sources, requirements, evidence, findings, actions, reviewers and outputs.
Decision authority	Retained by authorised human reviewers; the platform records and supports decisions but does not make statutory, clinical, regulatory or professional determinations.
Deployment posture	Can operate as a domain-specific application above existing systems of record, using files, structured data and approved integrations.
AI posture	Optional, bounded, schema-checked assistance that produces review proposals rather than silent project mutations.
Commercialisation posture	A reusable platform foundation that can be configured and validated with research and industry partners for multiple vertical applications.

2. Problem model and design principles

2.1 The recurring readiness problem

Across regulated, technical and operational domains, teams commonly face the same structural problem:

- the target milestone or decision is known, but the evidence required to support it is spread across documents, systems and people;
- requirements are interpreted in context and cannot always be reduced to a static checklist;
- evidence may be incomplete, stale, conflicting, low-confidence or relevant only to part of the case;
- automation can extract or suggest information but should not silently determine acceptance, compliance or readiness;
- unresolved matters are often buried in reports, email or meeting notes rather than converted into owned actions;
- handover between teams weakens the link between source evidence, professional judgement and the final readiness state.

© 2026 Axiom Systems Pty Ltd. All rights reserved.

2.2 Design principles

Principle	Specification consequence
Evidence before status	A readiness state should be explainable through linked evidence, review decisions and unresolved items.
Human authority retained	Automation and AI can prepare candidates and cues, but acceptance and controlled application remain explicit human actions.
Provenance is part of the data	Source, time, method, scope, version, quality flags and reviewer status remain attached to evidence and outputs.
Configuration over hard-coding	Target states, requirement sets, gates, roles, templates and adapters should be configurable for a domain and pilot.
Visible uncertainty	Unknown, conflicting, stale and low-confidence states are first-class conditions rather than silently coerced into pass/fail.
Actions close the loop	A gap or finding becomes a visible, assigned and monitored action with due state and closure evidence.
System-of-record respect	URE can reference and orchestrate authoritative data without pretending to replace every specialist source.
Deterministic core, optional intelligence	Core workflow transitions and reporting remain deterministic; AI support is optional and bounded.
Non-blocking long-running work	Extraction, external inspections and AI assistance run asynchronously so users can continue working while status remains visible.
Claim discipline	The platform distinguishes current capability, configured pilot work, proposed research and future commercial options.

3. Core operating model

1. Define target state
2. Capture context
3. Register evidence
4. Apply requirements and checks
5. Surface gaps/findings
6. Human review
7. Assign actions
8. Record readiness/outcome

> **Readiness is not a single score**
>
> A domain may use statuses, conditions, thresholds or maturity levels, but the authoritative result should remain connected to the evidence and decisions that support it. A score without explainability is treated as a presentation aid, not a substitute for review.

© 2026 Axiom Systems Pty Ltd. All rights reserved.

3.1 Inputs

- briefs and structured forms
- documents and files
- manual observations and reviewer notes
- approved public or enterprise data sources
- CAD/BIM and tabular artefacts
- optional machine-generated extraction or AI candidates

3.2 Controlled transformations

- normalisation into a case record
- source and version registration
- requirement applicability and evidence linking
- quality and confidence flags
- review proposal creation
- explicit acceptance/edit/rejection/defer decisions
- action generation and lifecycle tracking
- deterministic readiness evaluation and reporting

3.3 Outputs

- readiness summary and conditions
- evidence register and provenance trail
- gaps, conflicts and findings
- action plan with owners and status
- review and decision record
- controlled report fields and data interchange
- audit events and operational metrics

4. Logical architecture

Dimension	Description
Experience layer	Role-aware web interface, brief-first intake, domain selector, dashboards, evidence review, action boards and reporting.
API and application services	Authenticated project/case services, domain commands, validation, authorisation and integration endpoints.
Readiness orchestration	Target states, requirement applicability, evidence sufficiency, findings, gates, actions, transitions and deterministic outputs.
Configuration layer	Domain packs containing terminology, requirement templates, target states, roles, check logic, report schemas and source adapters.
Evidence and intake services	File registration, type-aware extraction, OCR where configured, tabular/CAD/BIM preparation, provenance and candidate creation.
Review proposal service	Queues machine- or user-generated candidates for

© 2026 Axiom Systems Pty Ltd. All rights reserved.

Dimension	Description
	accept/edit/reject/defer and explicit apply.
Async job service	Queued processing for extraction, inspections and AI assistance with queued/processing/ready/partial/failed/cancelled states.
Integration adapters	REST/API, files, CSV, events and approved domain sources; explicit warnings and manual fallback.
Persistence and audit	Projects/cases, evidence, decisions, actions, configuration versions, audit events and report artefacts.
Security and operations	Tenant/workspace isolation, least privilege, secret management, observability, backups and controlled support access.
Optional AI adapter	Provider-agnostic OpenAI-compatible interface configured through environment settings; structured output and claim checks.

4.1 Architectural boundary

URE is designed to sit between raw information and an authorised decision. It is neither a generic document repository nor an autonomous decision engine. Specialist systems remain authoritative for their own data; URE preserves references, status and warnings while creating a coherent readiness workflow around them.

5. Canonical information model

Entity	Purpose
Workspace / Tenant	Security and ownership boundary for organisations and projects.
Project / Case	The bounded subject of a readiness assessment.
Brief / Context	Purpose, scope, location, stage, constraints and known assumptions.
Target State	The milestone, condition or decision for which readiness is being assessed.
Requirement	A domain-defined condition, question, obligation or evidence expectation.
Check	Deterministic or reviewer-guided logic used to assess a requirement.
Evidence Item	A claim, observation, artefact or structured value offered in support of one or more requirements.
Source Reference	Origin, URI/file, version, timestamp, method, coverage and provenance information.
Quality / Confidence Flag	Currency, completeness, conflict, uncertainty, extraction confidence or reviewer caveat.

© 2026 Axiom Systems Pty Ltd. All rights reserved.

Entity	Purpose
Finding / Gap	An unresolved condition, inconsistency, deficiency, risk or information need.
Action	Owned next step with status, priority, due state, dependency and closure evidence.
Review Proposal	Candidate changes awaiting individual human disposition and explicit application.
Decision Record	Reviewer, authority, rationale, conditions, time and applied changes.
Readiness Assessment	Versioned outcome with supporting evidence, unresolved items and configuration version.
Integration Run / Job	Traceable execution of an extraction, source check, import or AI task.
Audit Event	Who did what, when, in which workspace, with relevant reason and result.

5.1 Relationship rules

- Evidence can support multiple requirements, but every link records the purpose and review status.
- A requirement may remain unresolved even when evidence exists if the evidence is insufficient, stale, conflicting or outside scope.
- Findings and actions retain links to the triggering requirement, evidence and review decision.
- A readiness assessment is versioned and records the configuration version used to calculate or compose it.
- Deleting or superseding an artefact does not erase historical decisions; the audit trail preserves the prior state and supersession relationship.

6. Workflow and state model

6.1 Reference case lifecycle

Reference state	Meaning
Draft	Case exists; target state and context are incomplete.
Contextualised	Target state, scope and known constraints are sufficiently defined for evidence planning.
Evidence assembling	Requirements are active; evidence, source cues and findings are being gathered.
Under review	The evidence set and unresolved matters are presented to authorised reviewers.
Ready	The agreed gate criteria are met with no unresolved blocking condition.
Ready with conditions	Progress is authorised subject to explicit conditions and owned

© 2026 Axiom Systems Pty Ltd. All rights reserved.

Reference state	Meaning
	actions.
Not ready	Blocking evidence gaps, findings or decisions remain.
Deferred / Cancelled	Assessment is paused or closed with rationale.

Configuration note. A domain pack may rename or refine states, but transitions remain typed, auditable and guarded. Invalid transitions are rejected rather than silently coerced.

6.2 Long-running operation lifecycle

Job state	Required UX/system behaviour
Queued	Request accepted and user can continue normal work.
Processing	Worker is active; status and available progress are visible.
Ready for review	Output is complete but has not been applied to the project.
Partial	Some outputs are available and limitations are explicit.
Failed	No silent fallback; error and retry/recovery path are visible.
Cancelled	User or system cancelled the operation with audit record.

Non-blocking UX rule

Completion of OCR, extraction, source inspection or AI assistance must not automatically steal focus or reposition the user. A visible status surface may offer “Check result”; only the user’s action moves focus to the relevant review area.

7. Evidence intake, extraction and provenance

7.1 Supported intake pattern

- manual brief and structured-form entry;
- PDF, DOCX, XLSX, CSV, text/RTF and image registration;
- CAD/BIM registration for DWG, DXF and IFC, with bounded metadata or conversion paths where configured;
- approved API/source checks with explicit source status and warnings;
- batch import through controlled templates;
- optional OCR and machine-assisted extraction with page, file and runtime bounds.

7.2 Provenance record

Provenance field	Specification
Origin	Source system, file, URL, person, device or generated process.
Time	Observed, issued, accessed, imported and reviewed timestamps as applicable.

© 2026 Axiom Systems Pty Ltd. All rights reserved.

Provenance field	Specification
Version	Document/model/data version, checksum or immutable identifier where available.
Method	Manual entry, extraction, OCR, API adapter, import or AI-assisted proposal.
Coverage	Which site, asset, period, population, requirement or scope the evidence represents.
Quality	Completeness, currency, conflict, confidence and known limitations.
Authority	Who supplied, reviewed, accepted or superseded the item.
Use	Requirements, findings, actions and outputs that rely on the evidence.

7.3 Acceptance boundary

Extraction is not evidence acceptance. A parsed value, OCR result, model metadata item or source cue enters a review context. A reviewer may accept, edit, reject or defer the candidate. Applying accepted content to the case is a separate controlled step. Low-confidence and unknown outcomes remain visible rather than being converted into a positive finding.

8. Human review and bounded AI assistance

8.1 Review proposal pattern

9. A user or system requests assistance for a permitted task.
10. The service receives a bounded project snapshot and task-specific schema.
11. The service produces candidate items with supporting references, limitations and runtime metadata.
12. Schema and claim checks reject malformed or unsupported output.
13. Valid candidates enter a Review Proposal; the project remains unchanged.
14. An authorised reviewer accepts, edits, rejects or defers each item.
15. Only accepted or edited items are explicitly applied through a controlled command.
16. The audit records proposal, disposition and application without retaining or exposing hidden chain-of-thought.

8.2 AI non-claims

- AI does not approve a planning pathway, certify compliance, provide clinical authority or make a regulatory determination.
- AI output is not automatically treated as source evidence.
- The engine does not claim universal reliability across all documents, sources or domains.
- A model confidence statement is not substituted for evidence quality or reviewer authority.
- Provider and model configuration are operational choices; domain validation remains necessary after any change.

8.3 Provider abstraction

The reference AI integration should use environment-configured base URL, model and API credential values through an OpenAI-compatible interface. This permits local or hosted models to be substituted without changing the readiness-domain contract. Secrets are never embedded in prompts, source files or client code.

© 2026 Axiom Systems Pty Ltd. All rights reserved.

9. Configuration and verticalisation

9.1 Domain pack contents

Domain-pack component	Purpose
Vocabulary	Domain labels, definitions, status language and user guidance.
Target states	Milestones and allowed readiness outcomes.
Requirement catalogue	Questions, obligations, evidence expectations, applicability rules and dependencies.
Evidence profiles	Expected artefact types, fields, provenance and quality criteria.
Roles and authorities	Who can contribute, review, approve, defer and apply changes.
Checks and gates	Deterministic checks, reviewer prompts, blocker rules and transition guards.
Source adapters	Approved integrations, warnings, rate/error handling and manual fallback.
Report schemas	Readiness summaries, decision records and controlled exchange formats.
Pilot metrics	Baseline definitions, outcome measures, telemetry and evaluation instruments.
Safety boundary	Prohibited claims, decisions and automated actions for the domain.

9.2 Reuse boundary

Reuse applies to the evidence, review, action and readiness orchestration architecture. Domain semantics, evidence sufficiency, professional roles, integrations and validation are not assumed to transfer automatically. Each vertical requires co-design and explicit verification by appropriate subject-matter experts.

10. Integration model

Integration class	Specification
Files and records	Upload/registration with hashes, versions, provenance and preview/extraction where available.
CSV/tabular exchange	Controlled import/export templates with validation and row-level errors.
REST APIs	Authenticated adapters with source identity, request/response status, timestamp and warning capture.
Events/webhooks	Optional event-driven updates with idempotency and replay

© 2026 Axiom Systems Pty Ltd. All rights reserved.

Integration class	Specification
	safety.
Identity	Organisation/workspace identity, invitation and role mapping; enterprise SSO is a configurable deployment extension.
Reports	PDF/Markdown/JSON or domain-specific templates generated from the same versioned evidence record.
External system links	Deep links and identifiers preserve the authoritative-system relationship without copying unnecessary data.

10.1 Adapter contract

- declare source and adapter version
- return success, partial, unavailable or failed status explicitly
- preserve timestamps and identifiers
- carry coverage and warning metadata
- avoid asserting authority beyond the source contract
- support retry/idempotency where applicable
- provide manual fallback and user-visible limitations

11. Security, privacy and audit boundary

Control area	Current/design boundary
Workspace isolation	Projects and evidence are scoped to an owner/organisation boundary; cross-owner access is not implied.
Least privilege	Commands and data access are role-checked; unavailable commands are disabled rather than merely hidden.
Controlled support	Owner-approved read-only or temporary-write support with reason, typed confirmation, expiry, revocation and audit.
Auditability	Material reads/writes, review dispositions, state transitions and support events are recorded.
Secret handling	Credentials are held in environment/secret stores and excluded from client bundles, prompts and reports.
Data minimisation	Collect only data necessary for the configured readiness purpose; preserve links to authoritative sources where copying is unnecessary.
Retention and deletion	Deployment-specific policies define retention, legal hold, export and deletion; historical decisions require controlled supersession.
AI data controls	Model routing, retention settings and permitted data classes are explicit configuration decisions; sensitive data use requires partner-specific approval.
No certification claim	This specification describes controls and design intentions; it does not

© 2026 Axiom Systems Pty Ltd. All rights reserved.

Control area	Current/design boundary
	claim ISO, SOC 2, IRAP or other certification.

12. Non-functional requirements

Quality attribute	Requirement
Traceability	Every reported status can be traced to configuration version, evidence, findings, decisions and actions.
Determinism	Core state transitions and readiness outputs are reproducible for the same accepted inputs and configuration.
Availability	Pilot deployment targets are agreed per partner; long-running tasks do not block normal UI work.
Performance	Interactive commands should return promptly; intensive work is queued and status-driven.
Resilience	Jobs are idempotent where practical, partial outcomes are explicit and retries do not duplicate accepted evidence.
Accessibility	Web UI targets WCAG-aligned keyboard, contrast, focus and semantic interaction patterns.
Observability	Health, job state, errors, integration status and usage are observable without exposing secrets.
Portability	Configuration and adapters are separated from domain-independent orchestration; deployment can move across suitable hosting environments.
Testability	Business rules, transitions, adapters and live/mock parity are testable through automated suites and domain acceptance cases.
Maintainability	Typed contracts, migrations, explicit configuration versions and bounded modules reduce hidden coupling.

13. Pilot deployment model

13.1 Three-stage pathway

Stage	Purpose
Stage A - Discovery and configuration	Define target decision, users, evidence map, current workflow, metrics, safety boundary and candidate domain pack.
Stage B - Research prototype	Configure the engine, use representative or controlled data, run expert walkthroughs and refine requirements and measures.

© 2026 Axiom Systems Pty Ltd. All rights reserved.

Stage	Purpose
Stage C - Controlled pilot	Invite bounded users/cases, operate support and audit controls, compare against baseline and evaluate usability, completeness, efficiency and decision quality.

13.2 Minimum pilot definition

- one clearly bounded target state
- one named user cohort and authorised reviewer role
- an agreed evidence set and source boundary
- a baseline workflow and measurement protocol
- defined non-claims and escalation rules
- representative cases
- success and guardrail metrics
- data/IP/publication arrangements
- exit, continuation and handover criteria

14. Demonstrated maturity and claim boundary

The URE concept is grounded in an existing vertical implementation, AxiomReady, rather than a greenfield architecture proposal. Internal evidence records describe a cloud-deployed, invitation-controlled planning-readiness prototype with brief/document intake, site/context cues, requirements and evidence orchestration, findings/actions, professional-review gates, deterministic reporting and review-gated AI assistance. The recorded clean technical baseline passed 204 of 204 backend tests in August 2026.

> **Critical maturity distinction**
>
> The TRL 6 evidence applies to the bounded AxiomReady planning-readiness implementation. It does not mean that proposed cyber, agriculture, health, energy or other verticals are already TRL 6. Those are candidate research applications to be co-designed, configured and validated.

Internal evidence basis

ID	Internal record	Date/baseline	Use in this pack
E1	Axiom Universal Readiness Engine Product Brief	August 2026	Defines the URE operating model, value proposition and human-decision boundary.
E2	AxiomReady - HIN0027 TRL 6 and Prototype Evidence	12 August 2026 baseline	Records a functional cloud-deployed planning-readiness vertical, its bounded workflow and 204/204 passing backend tests.
E3	AxiomReady Deployment and Pilot Readiness Statement	15 August 2026	Records deployment, private workspaces, controlled support and current pre-commercial status.

© 2026 Axiom Systems Pty Ltd. All rights reserved.

Boundary. The demonstrated maturity claim applies to the existing AxiomReady planning-readiness implementation and its bounded tested workflow. Other domain implementations in this package are prospective configurations requiring co-design, domain verification and pilot validation.

15. Candidate domain implementations

The following catalogue illustrates how the same orchestration pattern could be configured. It is an opportunity map, not a claim that every vertical is already implemented.

Candidate vertical	Illustrative target state	Research opening
Planning & Development	A defined planning or development milestone is supported by current, attributable evidence, with unresolved matters visible and owned.	How can evidence provenance, structured review and bounded AI assistance reduce repeated due-diligence work without displacing professional planning judgement?
Construction Compliance & Digital Engineering	The evidence required for a nominated design, certification, construction or handover gate is assembled, reviewed and traceable.	Can a configurable evidence graph improve traceability across BIM, documents and human review while reducing manual compliance reconciliation?
Cybersecurity & Operational Resilience	A selected cyber-resilience baseline is supported by reviewed control evidence, explicit exceptions and owned remediation.	How should evidence quality and reviewer confidence be represented so a cyber-readiness status is useful without becoming a misleading compliance score?
AI Governance & Responsible AI	An AI use case is demonstrably ready for an agreed lifecycle gate, subject to authorised human approval and documented residual risk.	Can a human-review-gated evidence workflow operationalise AI governance frameworks while preserving context-specific judgement and avoiding automated self-certification?
Agriculture & Farm Operations	A defined operational decision or seasonal plan is supported by appropriate evidence, reviewed assumptions and assigned next actions.	Can an evidence-and-review layer improve adoption of digital agriculture by making uncertainty, provenance and action ownership explicit rather than merely adding another data dashboard?
Biosecurity & Traceability	The enterprise or facility can demonstrate readiness for the agreed biosecurity scenario, including evidence, roles, triggers and response actions.	How can a configurable readiness model support proportionate biosecurity assurance across heterogeneous enterprises without imposing a one-size-fits-all checklist?
Energy Transition & Project Assurance	A nominated project gate is supported by current cross-discipline evidence, accepted dependencies and visible residual risks.	Can evidence-centred readiness orchestration improve cross-discipline project assurance and reduce late discovery of unresolved dependencies?
Environmental Compliance & Natural Capital	The organisation can demonstrate that relevant environmental obligations are understood, evidenced, reviewed and actively managed.	How can provenance-aware readiness workflows improve confidence in environmental evidence while accommodating evolving obligations and site-specific professional judgement?
Manufacturing & Quality Systems	A product, process or production run is ready for the agreed release gate with reviewed quality evidence and controlled deviations.	Can configurable evidence sufficiency rules and human review reduce release delays while maintaining or improving quality assurance?
Infrastructure & Asset Readiness	An asset, facility or operating transition is ready for service, handover, maintenance intervention or resilience testing.	Can a readiness layer across existing asset systems improve handover quality and operational resilience without duplicating systems of record?

© 2026 Axiom Systems Pty Ltd. All rights reserved.

Candidate vertical	Illustrative target state	Research opening
Research Commercialisation & TRL Progression	A research outcome or venture is ready for a defined translation, pilot, investment or deployment decision with evidence and assumptions visible.	Can evidence-linked readiness improve the quality and transparency of research translation decisions compared with stage-gate checklists alone?
Procurement & Supplier Assurance	A procurement or supplier decision is supported by current, reviewed evidence, transparent evaluation and owned conditions.	How can evidence provenance and review controls improve supplier assurance while limiting automation bias in evaluation?
Health Service Innovation	A service innovation is ready for a defined pilot or implementation gate, with clinical authority retained by authorised professionals.	Can a human-controlled readiness workflow improve implementation transparency and reduce missed dependencies in complex service change?
Emergency & Community Resilience	A community, organisation or facility can demonstrate readiness for an agreed hazard or continuity scenario.	How can readiness evidence remain current between exercises and incidents, and how should uncertainty be communicated to decision makers?

16. Research and validation agenda

Research hypothesis	Testable proposition
H1 - Configuration transfer	A common evidence/review/action model can support materially different domains while domain semantics remain independently configured.
H2 - Traceability	Linking status to evidence and review decisions improves explainability and reduces time spent reconstructing rationale.
H3 - Human gating	Explicit review proposals reduce automation bias and erroneous project mutation compared with direct AI write-back.
H4 - Readiness quality	Visible gaps, uncertainty and action ownership improve the quality of readiness decisions compared with document-only or checklist-only practice.
H5 - Workflow efficiency	A configured engine can reduce repeated information retrieval and re-entry without reducing professional review quality.
H6 - Cross-team continuity	A versioned evidence trail improves handover and consistency across roles and organisational boundaries.
H7 - Adoption	A brief-first, non-blocking interface and domain language improve usability relative to tool-centred data intake.

16.1 Candidate evaluation measures

- elapsed and active preparation time
- number of manual lookups and repeated data-entry steps
- evidence completeness and provenance coverage
- gaps or conflicts identified before the decision gate
- reviewer agreement and confidence

© 2026 Axiom Systems Pty Ltd. All rights reserved.

- false acceptance / false omission of blocking matters
- action ownership and closure
- rework after review
- user trust, workload and adoption intent
- system reliability and job completion

17. Limitations and non-claims

- URE is not a substitute for a regulator, certifier, clinician, engineer, planner, auditor or other authorised professional.
- The platform does not make a universal claim of compliance or readiness across all jurisdictions and contexts.
- Domain configurations require expert validation, current source definitions and controlled change management.
- Public and external sources can be unavailable, incomplete or outdated; warnings and manual fallback remain necessary.
- AI and extraction performance varies by data quality, model, language, document type and task; outputs remain review candidates.
- Security, privacy, retention and hosting requirements must be agreed for each partner and deployment.
- Research outcomes, efficiency gains and commercial demand must be measured; they are not assumed from the platform architecture alone.

18. Glossary and evidence basis

Term	Meaning in this specification
Readiness	The evidence-backed status of a defined target state, including conditions, uncertainties and unresolved matters.
Target state	The milestone, decision, operational condition or assurance objective being assessed.
Evidence	An attributable artefact, observation, value or reviewed claim used to support or challenge a requirement.
Requirement	A condition, obligation, question or evidence expectation relevant to the target state.
Finding	A gap, conflict, exception, risk or issue requiring review or action.
Review Proposal	A set of candidate updates that cannot change the case until individually reviewed and explicitly applied.
Domain pack	Versioned configuration of vocabulary, requirements, evidence profiles, roles, checks, sources, gates and reports.
Provenance	Information describing the origin, method, time, version, coverage and authority of data or evidence.
System of record	The authoritative external system for a particular data class or transaction.

© 2026 Axiom Systems Pty Ltd. All rights reserved.

Internal evidence basis

ID	Internal record	Date/baseline	Use in this pack
E1	Axiom Universal Readiness Engine Product Brief	August 2026	Defines the URE operating model, value proposition and human-decision boundary.
E2	AxiomReady - HIN0027 TRL 6 and Prototype Evidence	12 August 2026 baseline	Records a functional cloud-deployed planning-readiness vertical, its bounded workflow and 204/204 passing backend tests.
E3	AxiomReady Deployment and Pilot Readiness Statement	15 August 2026	Records deployment, private workspaces, controlled support and current pre-commercial status.

Boundary. The demonstrated maturity claim applies to the existing AxiomReady planning-readiness implementation and its bounded tested workflow. Other domain implementations in this package are prospective configurations requiring co-design, domain verification and pilot validation.

© 2026 Axiom Systems Pty Ltd. All rights reserved.