Universal Readiness Engine: A General Model for Evidence-Grounded, Review-Gated Readiness Systems

Adi Saric

Independent Researcher | Sydney, Australia

ORCID: 0009-0003-5963-100X | Correspondence: adi.saric@axiomsystems.info

Abstract

Projects, packages and cases are often declared ready through dispersed documents, workflow states, checklists and professional judgement. This paper proposes the Universal Readiness Engine (URE), a conceptual model for a narrower information-systems problem: representing whether an identifiable subject is sufficiently supported for a specified target under an explicit policy. URE treats readiness as a target-relative projection over requirements, evidence, findings, reviewed decisions and history, rather than as a document count, workflow position, maturity score or autonomous decision. It also distinguishes operational job state, candidate/informational state and authoritative domain state. Observations from OCR, source retrieval, rules or AI can inform proposals, but do not become authoritative semantic state merely by being generated or completed; a permitted, provenance-bearing, policy-governed semantic transition is required, with accountable human review wherever the applicable policy assigns authority to a reviewer.

The study uses a design-science-informed artefact analysis. Candidate constructs were extracted from a software-artefact family, abstracted away from planning-specific terminology, challenged against a narrower carbon-evidence implementation and a prospective farm-domain contract, and positioned against prior work in assurance cases, dynamic assurance, requirements traceability, provenance, goal-oriented requirements engineering, case and decision management, readiness frameworks, argumentation, continuous compliance and human-AI governance. The strongest implementation evidence is one planning-readiness reference implementation. It demonstrates operational feasibility of the proposed separation; it does not establish superiority, empirical effectiveness, or universality.

The contribution is therefore not any individual mechanism. Evidence-linked claims, provenance, traceability, review, rule evaluation, dynamic assurance and readiness assessment all have substantial antecedents. Subject to further formal and empirical validation, URE's proposed contribution is their operational synthesis around target-relative readiness semantics and an explicit candidate-to-

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

authoritative state boundary. “Universal” denotes an intended domain-independent model class, not a claim that every decision problem fits the model.

Keywords: operational readiness; assurance cases; evidence provenance; requirements traceability; case management; decision management; design science research; human-AI governance; controlled state transition

1. Introduction

Whether a submission is ready to lodge, a handover package is ready to transfer, a carbon evidence pack is ready for review, or a case is ready to progress looks like a simple status question. It is usually not. The judgement depends on an intended target, conditions that apply to that target, changing documents and records, missing or conflicting material, specialist review, and decisions made by actors with different authority. A green status, a completed task, or a workflow stage can be useful, but none alone explains why a conclusion is justified, what could invalidate it, or who was permitted to make it.

This paper asks whether that recurring problem class admits a domain-independent conceptualisation. It proposes the **Universal Readiness Engine (URE)**: a reference model for making a readiness conclusion visible, inspectable and operationally governable when it depends on connected requirements, evidence, findings and reviewed decisions. The model is deliberately not a universal decision-maker, legal interpreter, safety case language, workflow engine, or AI system. It does not determine that a document is true, an authority will approve an application, or a professional decision is correct.

The central research question is:

> How can readiness be represented as a target-relative, evidence-grounded and policy-governed state without allowing computational observations or proposals to acquire semantic authority merely through their generation?

Three contributions are proposed, with claim strength intentionally constrained by the evidence base.

1. **Target-relative operational readiness.** A conceptualisation of readiness as a policy-relative projection over the current authoritative state for a subject and target, rather than a scalar maturity level or generic workflow position.
2. **A governed candidate-to-authoritative transition.** A governance model in which computational and external outputs occupy a candidate state until a provenance-bearing, policy-authorised semantic transition makes a domain consequence authoritative, requiring accountable human review wherever the applicable domain policy requires it.
3. **A domain-independent readiness metamodel and reference architecture.** A small set of relations among context, target, conditions, evidence, findings, proposals, decisions, history and reporting, abstracted from an artefact family and challenged across domains of unequal maturity.

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

These are not priority claims over the constituent mechanisms. Assurance cases already organise claims, argument and evidence; dynamic assurance already treats evolving evidence; traceability and provenance already model relations and derivation; CMMN and DMN already support case and decision modelling; readiness frameworks already use evidence; and human-AI research already motivates appropriate oversight. The proposed contribution lies, if it survives subsequent evaluation, in a particular operational synthesis and semantic emphasis.

2. Research Method and Scope

2.1 Design-science-informed artefact analysis

The study is positioned as a **design-science-informed artefact analysis**, not as a controlled experiment or a retrospective claim that the whole development programme followed a formal Design Science Research Methodology (DSRM). Design science treats designed artefacts as vehicles for knowledge about a problem and solution space [1]. DSRM supplies a commonly used process vocabulary—problem identification, objective definition, design and development, demonstration, evaluation and communication [2]. Those sources support the stance that an implemented artefact can motivate conceptual knowledge; they do not turn an implementation review into outcome evidence.

The practical problem class was uncertainty at decision-relevant milestones: documents exist, but their relevance, sufficiency, currency, conflicts and required review are difficult to reconstruct. The artefact family embodied candidate constructs. This study then performed six bounded analytical activities.

4. **Repository analysis.** Types, persistence operations, contracts, tests and domain modules were inspected to identify explicit state, relationships and mutation rules.
5. **Construct extraction.** Candidate constructs were identified where the implementation retained stable identities and relations: targets, requirements, evidence, findings, proposals, decisions, actions and histories.
6. **Abstraction.** Planning-specific terms, interfaces and integration details were removed unless they expressed a reusable relation.
7. **Cross-domain challenge.** The planning implementation, a carbon-evidence implementation, and a farm-domain foundation contract were compared according to their actual maturity rather than treated as equivalent deployments.

8. **Falsification.** Candidate claims were challenged by asking whether they were already explained by assurance cases, workflow/case/decision standards, traceability, readiness models, argumentation, compliance automation or AI governance.

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

9. **Literature positioning.** Primary standards and peer-reviewed work were used to narrow claims and identify the remaining contribution question.

This method produces an artefact-derived conceptual proposal. It does not measure user outcomes, reviewer quality, safety, accuracy, economic value, adoption, or inter-rater reliability. It also does not prove a formal semantics, despite using lightweight notation.

2.2 Evidence hierarchy and transfer discipline

The planning-readiness implementation is the strongest evidence. It contains typed targets, requirements, evidence states, findings, review proposals, explicit apply records, lifecycle recalculation and regression tests. A carbon-evidence module is a **partial instantiation**: it has deterministic evidence-gap and calculation logic but not the full target/review/proposal model. The farm-domain material is a **prospective architecture contract**, not a mature implementation. This hierarchy supports a portability hypothesis, not a demonstration of universality.

The analysis also distinguishes three things often collapsed in product accounts: (a) an implementation property, (b) a proposed design principle, and (c) a theoretical contribution. For example, a specific reviewed-apply function is implementation evidence; the general separation between candidate and authoritative state is a design principle; its possible distinctiveness as part of URE remains a theoretical question.

2.3 Claim discipline

The following claim discipline governs the manuscript. A **repository implementation property** is a proposition that can be checked in code, tests or contracts. A **design principle** is an abstraction that appears useful for building a readiness system, but may have extensive antecedents. A **theoretical contribution** is a claim about a model's conceptual distinction relative to existing scholarship; it cannot be established merely because a codebase embodies it. An **empirical claim** concerns observed outcomes in use and requires a study design appropriate to those outcomes. These levels are not interchangeable.

This distinction matters because a readiness system can be persuasive in demonstration while still having untested causal value. The reference implementation can show that a project store is able to retain a proposal, require an explicit apply operation, and recalculate a status after a change. It cannot, without independent evidence, show that users understand the model, make better decisions, experience less automation bias, or achieve better approval outcomes. Likewise, the presence of a formal-looking projection does not make the model formally verified. The notation is used to identify assumptions that later work can test, not to conceal them.

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

The prior-art analysis was deliberately adversarial. Each candidate contribution was checked first for an established equivalent. Evidence-linked argumentation is addressed by assurance cases; artefact relations by traceability; lineage by provenance; state progression by workflow and case-management standards; deterministic logic by decision management; maturity and evidence accumulation by readiness frameworks; and human oversight by HCI and governance. Consequently, this paper treats URE as a **model of composition and operational semantics**. It does not claim ownership of the components from which it is composed.

3. Conceptual Model

3.1 Readiness is target-relative

Let p be a subject or case, t a target state, and τ a time. Examples of t include *submission-ready*, *ready for professional review*, *handover-ready* and *evidence-pack-ready*. A target is not merely a label for a workflow stage. It determines which conditions apply, what evidence is relevant, which findings block progress, what currency is required and which actor can make a decisive disposition.

Let the authoritative state at time τ be:

```
Aτ(p) = <C, T, Q, E, F, D, H>
```

where C is context, T targets, Q requirements/conditions, E evidence and explicit absence states, F findings/gaps/conflicts, D reviewed human decisions and H provenance/history. Let Πt be the policy applicable to target t: applicability, required dispositions, permitted evidence states, blocker rules, freshness conditions, actor authority rules, permitted automatic semantic transitions, human-review requirements and supersession rules. Readiness is the projection:

```
S(p, t, τ) = project(Πt, Aτ(p))
```

The projection should normally be a vector, not an undifferentiated score:

```
S = <eligibility, coverage, blockers, uncertainty,
  review-status, source-currency, explanation-anchors>
```

A scalar or traffic light may be rendered from S, but it is a presentation of policy-defined state rather than the theory's primary object. A Boolean form is possible only after policy has specified its terms:

```
Ready(p,t,τ) iff eligible(t,p,τ)
    ∧ complete(Q_t,E_t,D_t,Πt)
     ∧ noBlocking(F_t,D_t,Πt)
   ∧ reviewedIfRequired(D_t,Πt)
    ∧ currentIfRequired(H_t,Πt)
```

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

This formulation does not assert that every readiness judgement should be Boolean, deterministic, or machine-computable. It makes explicit the information that a defensible computation or human assessment must be able to inspect.

3.2 Informational, operational and authoritative state

URE separates three state types because they have different meanings and permissions.

State type	Typical examples	May change readiness by itself?	Purpose
Operational job state	queued, processing, failed, completed	No	Track work performed by a system or service
Candidate/informational state I_τ	OCR output, source refresh, AI suggestion, unresolved conflict, proposal	No	Preserve observations and alternatives for assessment
Authoritative domain state A_τ	accepted evidence, requirement disposition, confirmed finding, completed qualified review	Yes, through policy	Support the readiness projection

The distinction is not an assertion that candidate data are useless or false. They can be highly reliable, and a domain policy may allow deterministic automation in carefully bounded conditions. The claim is semantic: a job completing is evidence that a job completed, not by itself evidence that an extracted proposition has been accepted for a target. Accordingly:

$$\text{job.completed} \neq \text{evidence.accepted}$$
$$\text{job.completed} \neq \text{requirement.satisfied}$$
$$\text{job.completed} \neq \text{target.ready}$$
$$\text{AI.response.generated} \neq \text{decision.authorised}$$

3.3 Controlled semantic mutation boundary

The candidate layer may contain observations, extraction results, rule output, inferred risks, external source records and proposals. A candidate does not become authoritative because it was generated, retrieved, inferred, or because its job completed. URE therefore defines a **controlled semantic mutation boundary** governed by the applicable target policy.

$$g \in G\Pi_t(A_\tau, o)$$
$$\mu\Pi_t : (A_\tau, o, g) \to A_{\tau+1}$$

Here, $G\Pi_t(A_\tau, o)$ is the set of authorisations valid for proposal or input o in the current authoritative state under Π_t. The authorisation g records the policy basis and transition class. It may be g_h, an accountable human-reviewed authorisation linked to a decision in D, when policy requires human authority. It may be g_a, a bounded deterministic authorisation, only when policy explicitly permits that exact transition class and its preconditions. The boundary is crossed only when policy conditions are

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

met, relevant references are current, conflicts requiring resolution are handled, and the resulting mutation is recorded in H. The formalism is descriptive, not a proof of correctness.

"Semantic write barrier" is retained as a short explanatory label for this pattern, but **controlled semantic mutation boundary** is the normative term in this paper. Its governance concept is the **governed authority boundary**. These terms avoid implying a particular database implementation or an approval button. They specify the class change from candidate information to authoritative domain state, the policy basis and authority that permit it, and the history needed to contest it. Human review remains essential whenever policy assigns the authority to a human, but an authorised deterministic transition does not become an autonomous decision architecture. Nor does any recorded review guarantee meaningful oversight; it makes the authority path explicit and auditable.

3.4 Minimum metamodel and relations

The reference metamodel uses a small set of connected objects rather than a large universal taxonomy.

Construct	Role in the model	Essential relations
Subject/case and context	Bounds the unit of assessment	has target, constraints, participants
Target and policy	Defines the intended milestone and applicable semantics	activates requirements, rules and authority conditions
Requirement/condition	States what bears on a target	applies to target; is supported, unresolved or superseded
Evidence or explicit absence	Provides potential support, uncertainty or contradiction	concerns requirement; has provenance, status and currency
Finding/gap/conflict	Makes insufficiency or disagreement first-class	affects target/requirement/evidence; may block
Proposal	Holds candidate change without authority	derives from observation; is reviewed or assessed against a permitted automatic transition
Decision and authority	Records human-reviewed disposition and who made it	reviews proposal/condition; provides g_h where policy requires
History/provenance	Explains lineage, authority basis and prior state	records source, activity, actor/service, policy version, time and resulting mutation
Action/communication	Coordinates resolution and reporting	addresses a finding; does not itself prove readiness

Target applicability is a first-class relation, `applies(q,t)`, because an item can be relevant to one milestone but not another. Explicit absence is likewise a state, not just a null: *not applicable*, *missing*, *requested*, *under review*, *rejected*, and *superseded* have materially different consequences.

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

3.5 Principles

The earlier principles survive with narrower classification.

- **Target-state principle (design principle):** readiness is always relative to a stated target and policy.
- **Evidence-relative principle (inherited concept, URE application):** artefact existence does not establish target support.
- **Explicit-gap principle (design principle):** absence, uncertainty, conflict and staleness should be modelled as inspectable state.
- **Governed-authority-boundary principle (proposed architectural synthesis):** candidate information is not authoritative merely because it was computed or retrieved; semantic authority requires a transition permitted by policy, with human review where policy requires it.
- **Controlled non-monotonicity (inherited concept, URE application):** later evidence, a reopened finding or a changed source may legitimately supersede a conclusion; dynamic assurance is clear antecedent work [6].
- **Provenance-and-contestability principle (inherited concept, URE application):** lineage supports inspection, not truth, correctness or immutable audit [11].
- **Domain-separation principle (reference-architecture principle):** an invariant structure should be separated from domain rule packs, terms and policies.
- **Assistance-without-authority principle (proposed architectural synthesis):** AI and other assistance may help generate candidate state while authority remains governed by an explicit policy-authorised transition; generic AI output receives no semantic authority merely because confidence is high.

3.6 Proposed URE-applicability class

URE is proposed for problems that have all or most of the following: an identifiable subject; at least one target; target-relevant conditions; evidence or explicit absence bearing on those conditions; representable gaps or conflicts; identifiable authority for state-changing decisions; an explanation need; and material temporal change. This is an applicability test, not a theorem. Problems whose answer is an immediate computation with no continuing evidence, authority or explanatory burden may not need URE.

3.7 Policy, explanation and controlled non-monotonicity

The policy object deserves explicit treatment. Πt is not merely a set of rules encoded in a table. At minimum it identifies: the target's scope and applicability relation; which conditions are mandatory, optional or risk-acceptable; what evidence states may contribute to coverage; what is blocking; what source currency means; which decision types and roles are sufficient; which transition classes are explicitly permitted to mutate semantic state automatically; which transitions require accountable

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

human review; and the conditions under which a prior review is superseded. A domain may express these elements in prose, a rule pack, DMN, a database schema or an assurance process. URE does not prescribe the representation.

This policy-dependence explains why the model cannot infer truth from evidence. A source may be authentic yet out of scope; a document may be current but insufficient; a condition may be complete only after a qualified reviewer records a disposition. Equally, an accepted-with-risk item may be compatible with a target in one policy but block another. The readiness function therefore evaluates a domain's declared semantics rather than discovering them from documents.

The model is intentionally non-monotonic in an operational sense. Let A_τ support a ready projection. A later source refresh, conflicting report, changed target policy or reopened finding can result in $A_{\tau+1}$ whose projection is no longer ready. The important qualification is materiality. Not every new observation should invalidate readiness. A candidate observation becomes relevant only through the policy and, where policy requires it, the controlled semantic mutation boundary. The result avoids two symmetric errors: immutable dashboards that ignore later evidence, and volatile dashboards whose status changes whenever an unreviewed model produces text.

An explanation is consequently more than a list of documents. For any displayed component of $S(p,t,\tau)$, explanation anchors should support questions such as: Which target and policy version were used? Which conditions were applicable? Which evidence and absence states mattered? Which findings block or qualify the result? Which human decision, or which bounded automatic authority basis, permitted the mutation? What actor or service executed it, when, against which prior state, and with what resulting state? These questions are compatible with provenance standards but extend lineage into a policy-relative operational explanation.

4. Reference Architecture

The architecture separates domain semantics from invariant relations. A planning rule pack may contain statutory triggers and consultant roles; a carbon pack may contain reporting boundaries and accepted evidence types; a handover pack may contain warranties and completion criteria. None should redefine the core distinction between candidate and authoritative state.

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

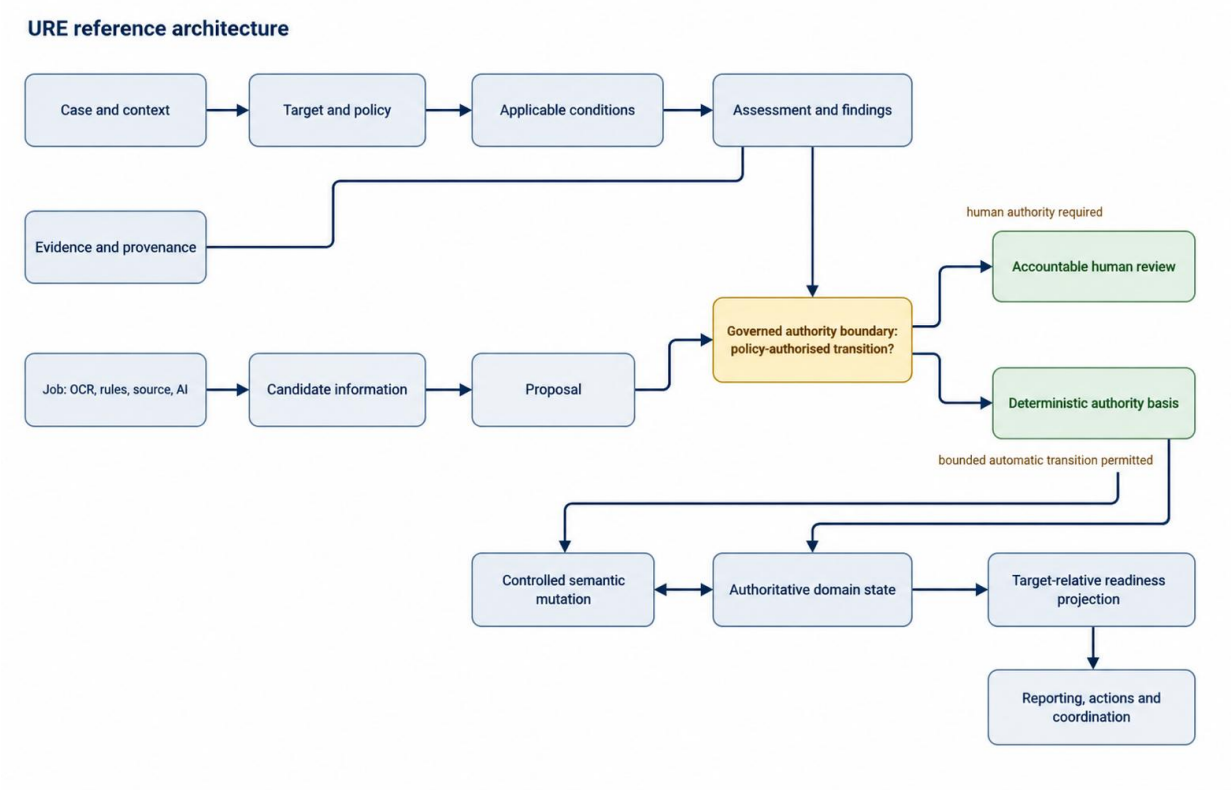


Figure 1. URE reference architecture.

The reference implementation demonstrates this pattern in bounded form. Its target entities connect requirements, findings and actions. Evidence states distinguish accepted, missing, requested and review-required material. Review proposals retain source and item decisions; the reviewed apply operation skips unresolved conflicts, non-accepted items and previously applied changes, then records what was applied or skipped. Its asynchronous OCR, extraction, source and AI contracts expressly avoid treating job completion as evidence acceptance, authority confirmation or readiness change. A lifecycle calculation can supersede a completed professional review when a blocker, warning or source change invalidates its conditions. These are implementation properties, not claims of a formal proof or a fully general engine.

The same implementation distinguishes internal submission readiness from a recorded external submission and from an authority outcome. This is important: being ready to submit is not evidence that submission occurred, and submission is not evidence of external approval.

4.1 Operational consequences of the architecture

The architecture has consequences for interface and service design. A system should avoid presenting a completed extraction, an accepted candidate field, a finished job and a ready target as visually indistinguishable states. They have different epistemic and governance meanings. At minimum, a user

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

should be able to identify the state class, the source, the confidence where relevant, the current reviewer disposition, and whether the item contributes to the readiness projection.

The approach also constrains automation. A source-refresh service may legitimately create a candidate notice that a public record has changed. A rule engine may identify that a required field is absent. An AI service may suggest a linked action or a possible conflict. Whether these outputs may automatically create a finding, modify a condition, accept evidence, close a blocker, or change readiness is a policy question. A defensible design makes the permitted transition class, preconditions and authority basis explicit rather than relying on the accidental order of events in a queue or API route.

This is not an argument against automation. A policy can define narrow automatic mutations when their inputs, semantics and failure behaviour are sufficiently controlled. For instance, an integrity-checked import may automatically register a document as received without treating it as accepted evidence; a deterministic calculation may update an estimate while leaving the relevant evidence status unreviewed. Such a mutation must record its policy/version, authority basis, triggering input, actor or service, timestamp, prior state and result. The boundary is therefore a way to specify *which* automation acquires authority, under what preconditions and with what provenance. It is stricter than “human in the loop” but is not necessarily “human in every loop.”

Finally, the model positions reporting and coordination downstream of authoritative state. A missing-item email, consultant request or action card can be generated from a finding and remain traceable to it. This preserves a distinction that operational systems often lose: an action is a response to a readiness problem, whereas a policy-valid disposition is what may resolve the problem. A task management system can therefore be a useful URE component without becoming the readiness authority.

5. Related Work and Theoretical Positioning

5.1 Assurance cases, GSN and SACM

Assurance cases are the closest conceptual family. SACM is an OMG metamodel for structured assurance cases, including argumentation and evidence artefacts [3]. GSN gives a standard notation and guidance for engineering arguments [4]. They demonstrate that claim-argument-evidence structure, auditable evidence packages and incomplete argument representation are established. URE must not claim to invent them.

The proposed distinction is a difference in primary operational question, not an exclusivity claim. An assurance case asks why a claim about a system or property should be believed and represents structured argument. URE asks, for a current target and policy, what readiness state should be projected

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

from authoritative conditions, evidence, unresolved findings and decisions, and what permitted transitions could change it. A URE instance may contain or link an assurance case; an assurance case can also implement readiness-like behaviour. The overlap is substantial and is a central novelty risk.

The practical relationship to prescriptive regimes is also relevant. Hawkins et al. show that assurance cases and prescriptive software safety certification can be complementary rather than mutually exclusive [5]. The same caution applies here. A domain policy may contain prescriptive conditions, an assurance argument may explain their sufficiency, and a URE-style projection may make the current operational status of those conditions visible. None of the three cancels the need for the others.

5.2 Dynamic and continuous assurance

Denney, Pai and Habli's dynamic safety cases explicitly address through-life assurance and changing evidence [6]. The existence of dynamic evidence, revisable arguments and non-monotonic assessment is therefore established prior art. URE's more limited emphasis is not that new evidence can change a conclusion, but that an observation changes an *authoritative readiness projection* only if it becomes semantically material under policy and crosses the controlled mutation boundary. That distinction may be useful in systems where an extraction job or a model output otherwise looks like a state decision; it is not yet demonstrated to be absent from dynamic-assurance implementations.

5.3 Goal-oriented requirements engineering and traceability

Goal-oriented requirements engineering studies refinement from goals to requirements, responsibilities and obstacles [8]. Requirements traceability studies links across the life of software artefacts and records unresolved technical and organisational challenges [7]. Both are direct antecedents for target-to-condition relations and evidence lineage. URE adds neither goal decomposition nor traceability. Its proposed extension is to place these relations in a current target projection that includes evidence state, finding disposition, decision authority and an operational readiness outcome. Whether that integration has practical benefit is an empirical question.

5.4 Provenance

PROV-O supplies a domain-independent ontology for representing and exchanging information about entities, activities and agents involved in producing a thing [11]. URE is a consumer or possible domain specialisation of provenance concepts, not a replacement for them. Provenance tells a reviewer what generated or derived an artefact; it does not alone state whether that artefact satisfies a policy, which finding blocks a target, or which reviewed mutation is permitted. Conversely, a URE history does not make evidence authentic, complete, correct or immutable.

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

5.5 Workflow, case and decision management

Workflow patterns establish reusable control-flow concepts [9]. CMMN supports evolving cases, discretionary work, milestones, knowledge workers and case files [10]. DMN standardises decision requirements and decision logic [12]. URE may use all three. CMMN is especially close where work is knowledge-intensive and order is unpredictable; DMN is close where policy can be encoded deterministically.

The proposed difference is again one of emphasis. CMMN primarily models what can or should occur in a case; DMN primarily models what decision follows from specified input and logic. URE adds a defined target-relative state explanation: which conditions apply, what evidence and decisions are authoritative, what remains unresolved, and why the target is or is not ready. CMMN or DMN can be extended to do this, and URE can be implemented through them. It is not a claim of replacement.

5.6 Readiness and maturity models

Readiness is already a mature research and policy topic. The GAO Technology Readiness Assessment Guide evaluates technology maturity and whether it is sufficiently developed for system integration [13]. System Readiness Level work combines technology and integration readiness into a system-maturity index [14]. Hughes et al. propose a Service Readiness Level framework that aligns evidence and assurance/exit criteria to stages of digital-health scale-up [15]. These are decisive counterexamples to any claim that URE introduces evidence-based readiness.

URE's proposed difference is its unit of analysis: a particular subject relative to a particular operational target under a stated policy, represented as an inspectable graph rather than principally a level or composite maturity measure. This does not make URE categorically different: a readiness framework can include evidence, targets, review and dynamic state. The distinction should be tested empirically rather than asserted as a taxonomy boundary.

5.7 Computational argumentation

Dung formalised acceptability in abstract argumentation frameworks and showed its connection to non-monotonic reasoning [16]. AIF addresses interchange of arguments and their structures [17]. Those traditions provide possible formal tools for support, attack, conflict and acceptability. URE does not need to become an argumentation framework. It currently treats findings and conflicts as operational state and leaves the semantic calculation to domain policy. Future work may map selected URE relations to argumentation semantics where that produces clearer conflict handling or explanation.

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

5.8 Continuous compliance and evidence automation

Compliance-as-code research shows that control artefacts, policies and evaluation results can be represented and assessed continuously [18, 19]. This is a strong adjacent implementation family, not merely a vertical market. It already connects controls, evidence, evaluation and remediation. URE should therefore be understood, at best, as a proposed generalisation of the structure needed when *compliance* is one possible target semantics. A compliance implementation may already constitute a URE-like instance under the paper's definition. The remaining question is whether the target-relative policy and candidate-authority distinction travel usefully beyond compliance.

5.9 Human-AI oversight and decision authority

Human-AI interaction guidance and NIST's AI RMF establish broad expectations around explanation, oversight, risk management and trustworthy use [20, 21]. URE does not claim human oversight for AI as new. Its narrower architectural proposition is to model computational outputs explicitly as non-authoritative candidate state until a policy-governed mutation occurs. That proposition should be evaluated against automation bias, reviewer workload, calibrated reliance and actual authority practice. It is compatible with existing guidance, not a substitute for it.

5.10 Synthesis and comparison

The table deliberately records emphasis, not competitive feature superiority. "Possible" means an approach can support the capability in an implementation; it does not mean the literature excludes it.

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

Approach	Primary object	Evidence linkage	Target-relative readiness	Dynamic state	Governed authority boundary	Provenance	Controlled candidate → authoritative transition
Checklist	Listed items	Possible	Possible	Implementation-dependent	Possible	Possible	Not primary
Workflow/BPM	Control flow	Possible	Not primary	Central	Possible	Possible	Possible
CMMN	Evolving case and work	Supported	Possible	Central	Supported	Possible	Implementation-dependent
DMN	Decision logic	Input-centred	Possible	Possible	Implementation-dependent	Possible	Implementation-dependent
Requirements traceability	Artefact relations	Central	Possible	Supported	Not primary	Supported	Not primary
Assurance case/SACM	Claim, argument, evidence	Central	Possible	Supported	Possible	Supported	Possible
Dynamic assurance	Evolving assurance argument	Central	Possible	Central	Possible	Supported	Possible
Readiness/maturity model	Maturity/readiness assessment	Supported	Often stage/level-centred	Supported	Possible	Possible	Not primary
Continuous compliance	Controls, policies and assessed state	Central	Compliance-specific	Central	Implementation-dependent	Supported	Possible
URE (proposed)	Current justified readiness for a target	Central	Central	Central	Central by design; human review where policy requires	Central for explanation	Proposed characteristic

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

Under the definition proposed here, a checklist becomes **URE-like** when it is connected to target policy, requirement applicability, evidence/absence state, findings, reviewed decisions, provenance and a readiness projection. This is a definition of the proposed model, not a claim that adjacent disciplines fail to contain those relations.

5.11 Interpretation of the comparison

The comparison table should not be read as a claim that URE has capabilities the other families lack. Many implementations of assurance cases, CMMN, DMN, continuous compliance or readiness assessment can represent target states, evidence, provenance and governed transitions. Conversely, a purported URE could degenerate into a checklist if its relations and policies are not actually implemented. The table identifies each family's primary conceptual centre and the areas where a URE design must reuse, map to, or be evaluated against established work.

This has two practical implications. First, an implementation should favour interoperability over reinvention. PROV-O can represent provenance; SACM can represent an assurance argument; DMN can hold decision logic; CMMN or workflow technology can coordinate work. URE need not create rival standards for each. Its reference architecture is best treated as a coordinating metamodel that states when these objects matter to a readiness target and when their effects become authoritative.

Second, the proposed controlled semantic mutation boundary should not be assumed to distinguish URE automatically. A mature compliance platform may already have an equivalent in its promotion controls, change-approval process and audit trail. A dynamic assurance environment may already distinguish monitored information from accepted argument elements. The research question is therefore comparative and operational: does making this boundary explicit in a general readiness model provide clearer explanation, safer automation and more controllable adaptation than leaving it implicit in adjacent tooling?

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

6. Cross-Domain Analysis

The model's invariant core can be expressed in several domain vocabularies without changing its proposed relations.

Domain	Target	Conditions/evidence	Typical findings	Authority boundary
Planning submission	Submission or professional-review ready	controls, drawings, reports, source snapshots	missing report, stale source, conflict	qualified review and external authority remain distinct
Carbon evidence	Evidence pack or accepted-work review ready	boundary, invoices, declarations, calculations	unsupported quantity, out-of-boundary evidence	accepted evidence depends on policy and review
Handover/tender	Handover or package ready	warranties, certificates, manuals, as-builts	missing certificate, incomplete item	package acceptance and contractual outcome are separate
Farm readiness (prospective)	Review-ready case	structured input and source candidates	missing/conflicting inputs	review proposal before canonical change

The planning implementation supports the core most comprehensively. CarbonWork supports evidence registration, source quality, deterministic gap detection and calculation but does not yet demonstrate full reviewed target-state semantics. FarmerReady documents a proposed flow from source/input through normalisation, review proposal and human decision to controlled canonical application; it should not be cited as an operational validation. The mappings establish only structural plausibility.

6.1 Worked conceptual scenario

Consider a package with two targets: *ready for professional review* and *ready for external submission*. A report is uploaded and an extraction job identifies a potential conflict with a required drawing. In J_τ, the job is complete. In I_τ, there is an extracted observation and an AI or rule-generated proposal. Neither fact changes the package's authoritative status. A reviewer may reject the proposal, defer it pending another source, or accept/edit it to create an authoritative finding. Only the permitted mutation adds the finding to $A_{\tau+1}$; the next projection can then show a blocker for the professional-review target. The external-submission target may remain ineligible because its professional-review condition is not complete.

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

Suppose the conflict is resolved and the reviewer records a policy-valid disposition against current evidence. The authoritative graph can now support readiness for professional review. That conclusion does not record a submission. A later external submission record is a separate authoritative fact. A subsequent authority response is separate again. This scenario does not depend on planning terminology: in carbon work, a supplier declaration may be the evidence; in handover, a warranty; in a farm case, a verified structured input. The target policy and actor roles vary, while the distinctions among job, candidate, authoritative state and external outcome remain intelligible.

6.2 Counter-scenarios and limits of transfer

Some domains may fail the proposed applicability test. A system that calculates a single sensor threshold from immutable calibrated data might need a decision rule and audit log but not a readiness metamodel. A fully formal verification system may produce a proof obligation whose authority is defined by its formal calculus, not a human review workflow. A high-volume routine process may permit deterministic automatic promotion under a carefully specified policy. URE should not be deployed merely because a process has documents and tasks. Its value hypothesis is strongest where heterogeneous evidence, temporal change, ambiguity and accountable human judgement make a target conclusion difficult to reconstruct.

Even within applicable domains, the generic core cannot resolve domain questions. It cannot decide which planning control applies, whether a carbon calculation complies with a reporting standard, whether a warranty is legally sufficient, or whether a source is authentic. Those questions belong to domain rule packs, professional practice and, where relevant, external authorities. URE can retain their reasoning and state consequences; it does not supply their substantive truth conditions.

7. What URE Does Not Claim

URE does not claim to replace domain expertise, assurance cases, workflow or case management, decision standards, provenance standards, or compliance automation. It does not establish legal, regulatory, professional or physical truth; authenticate evidence; eliminate human judgement; guarantee correct decisions; make an AI output authoritative; provide a universal maturity scale; or prove every domain fits the applicability class. Traceability supports explanation and auditability, not correctness. A human review event records an authority-bearing transition, not necessarily an adequate review.

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

7.1 Evaluation propositions

The model is valuable only if it produces testable consequences. The following are propositions for evaluation, not claimed results.

P1 — explanation completeness. Compared with a document repository or flat task list configured for the same milestone, an implementation of the model should enable a reviewer to retrieve the target, applicable conditions, evidence/absence states, blockers, reviewed decisions and lineage supporting a displayed conclusion with fewer unsupported assumptions. This can be evaluated through scenario-based explanation tasks, not through code inspection alone.

P2 — controlled automation. Where AI, extraction or source services are used, explicit candidate state and controlled mutation should reduce the number of accidental or unexplained authoritative changes compared with an architecture in which job results directly update business state. This requires fault injection, audit review and measurement of both false promotion and missed escalation. It may also reveal a cost: additional review latency.

P3 — calibrated reliance. By displaying a distinction between candidate and authoritative information, reviewers may be less likely to treat a high-confidence suggestion as an accepted case fact. This is a human-factors proposition. It cannot be inferred from the existence of the distinction; tests must assess automation bias, correction behaviour and understanding of authority.

P4 — transferable structure. Two independently built non-planning instantiations should be able to preserve the core relations while substituting their own policy and evidence vocabularies. Success would require more than matching labels: target applicability, explicit gaps, decision authority, history and readiness projection should operate coherently in the new domains.

P5 — interoperability rather than replacement. Mapping a URE instance to PROV-O, SACM, CMMN and DMN should expose which parts are duplicative, complementary or missing. A result showing that URE is a thin configuration of existing standards would narrow the theoretical claim but could still yield practical implementation guidance.

These propositions also identify trade-offs. Greater provenance and review control can create data-entry burden, reviewer queues and a misleading appearance of assurance. More automation can improve timeliness but blur semantic authority. A good evaluation should therefore measure both benefit and governance cost rather than assuming that a richer model is always better.

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

7.2 Design misuse risks

The model can be misused in predictable ways. A team may encode a policy that is wrong, stale or too coarse and then treat the resulting status as objective truth. An organisation may assign nominal reviewer roles without training, capacity or independence. A dashboard may conceal uncertainty behind a score. A proposal workflow may become a bureaucratic bottleneck, encouraging users to bypass it in untracked channels. An AI system may still anchor a reviewer's judgement even when its output is technically non-authoritative. Finally, provenance records can create a false sense of immutability or authenticity when they merely describe asserted lineage.

These risks imply that a mature implementation needs policy stewardship, role management, source-quality practices, transparent uncertainty, exception handling, audit procedures and periodic review of automation effects. They are organisational requirements around the architecture, not features the metamodel can supply by itself.

8. Limitations and Research Agenda

The most significant limitation is evidential: one implementation is mature, the others are not. The paper has no field study, comparison group, user evaluation, reliability analysis, or formal proof. "Universal" remains an architecture aspiration. The mutation-boundary proposal may already be present, under other names, in assurance, case-management or compliance systems; the present review does not establish novelty priority.

Future work should: (1) conduct a systematic literature review with defined databases and inclusion criteria; (2) formalise policy, authority and conflict semantics; (3) map URE constructs to SACM, PROV-O, CMMN and DMN; (4) test transfer into at least two fully implemented non-planning domains; (5) evaluate explanation quality, review burden, error recovery and automation bias; (6) compare target-projection outputs with existing readiness and assurance practice; and (7) examine governance failure modes, including superficial review, role ambiguity and stale policy.

9. Conclusion

URE is proposed as a restrained model for a recurrent operational question: whether a subject is currently ready for a particular target under an explicit policy, and why. It treats readiness as a projection over authoritative requirements, evidence, findings, decisions and history. Its central architectural pattern separates operational work and candidate information from authoritative state

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

through a controlled semantic mutation boundary. The pattern is demonstrated in a planning-readiness reference implementation and only partially or prospectively in two other domains.

The literature review materially narrows the claim. URE does not introduce evidence-based readiness, traceability, provenance, dynamic assurance, case management, rule evaluation, argumentation or human oversight. Its possible contribution is an operational synthesis: target-relative readiness semantics coupled with explicit governance over when information is permitted to change readiness. That contribution is plausible enough to warrant formal and empirical study, but not yet strong enough to claim universality, superiority or theoretical priority.

Declarations

Author and correspondence. Adi Saric, Independent Researcher, Sydney, Australia. Correspondence: adi.saric@axiomsystems.info. ORCID: 0009-0003-5963-100X.

Funding. This research received no external funding.

Conflict of interest. The author is the creator and maintainer of the Universal Readiness Engine concept and related implementation artifacts. This potential intellectual and commercial interest is disclosed; the paper's claims remain subject to the stated limitations and require independent validation.

Acknowledgements. None.

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

References

[1] Hevner, A. R., March, S. T., Park, J., & Ram, S. (2004). Design science in information systems research. *MIS Quarterly*, 28(1), 75-105. https://doi.org/10.2307/25148625

[2] Peffers, K., Tuunanen, T., Rothenberger, M. A., & Chatterjee, S. (2007). A design science research methodology for information systems research. *Journal of Management Information Systems*, 24(3), 45-77. https://doi.org/10.2753/MIS0742-1222240302

[3] Object Management Group. (2023). *Structured Assurance Case Metamodel (SACM), Version 2.3*. https://www.omg.org/spec/SACM/2.3/PDF

[4] Assurance Case Working Group. (2021). *Goal Structuring Notation Community Standard, Version 3*. Safety-Critical Systems Club. https://scsc.uk/r1386.pdf

[5] Hawkins, R., Habli, I., Kelly, T., & McDermid, J. (2013). Assurance cases and prescriptive software safety certification: A comparative study. *Safety Science*, 59, 55-71. https://doi.org/10.1016/j.ssci.2013.04.007

[6] Denney, E., Pai, G. J., & Habli, I. (2015). Dynamic safety cases for through-life safety assurance. In *2015 IEEE/ACM 37th IEEE International Conference on Software Engineering*, Vol. 2, 587-590. https://doi.org/10.1109/ICSE.2015.199

[7] Winkler, S., & von Pilgrim, J. (2010). A survey of traceability in requirements engineering and model-driven development. *Software and Systems Modeling*, 9, 529-565. https://doi.org/10.1007/s10270-009-0145-0

[8] van Lamsweerde, A. (2001). Goal-oriented requirements engineering: A guided tour. In *Proceedings of the Fifth IEEE International Symposium on Requirements Engineering*, 249-262. https://doi.org/10.1109/ISRE.2001.948567

[9] van der Aalst, W. M. P., ter Hofstede, A. H. M., Kiepuszewski, B., & Barros, A. P. (2003). Workflow patterns. *Distributed and Parallel Databases*, 14(1), 5-51. https://doi.org/10.1023/A:1022883727209

[10] Object Management Group. (2016). *Case Management Model and Notation (CMMN), Version 1.1*. https://www.omg.org/spec/CMMN/1.1/PDF

[11] Lebo, T., Sahoo, S., & McGuinness, D. (Eds.). (2013). *PROV-O: The PROV Ontology*. W3C Recommendation. https://www.w3.org/TR/prov-o/

[12] Object Management Group. (2024). *Decision Model and Notation (DMN), Version 1.5*. https://www.omg.org/spec/DMN/1.5/PDF

[13] U.S. Government Accountability Office. (2020). *Technology Readiness Assessment Guide: Best Practices for Evaluating the Readiness of Technology for Use in Acquisition Programs and Projects* (GAO-20-48G). https://www.gao.gov/products/gao-20-48g

[14] Sauser, B. J., Ramirez-Marquez, J. E., Henry, D., & DiMarzio, D. (2008). A system maturity index for the systems engineering life cycle. *International Journal of Industrial and Systems Engineering*, 3(6), 673-691. https://doi.org/10.1504/IJISE.2008.020680

[15] Hughes, J., Lennon, M., Rogerson, R., & Crooks, G. (2021). Scaling digital health innovation: Developing a new service readiness level framework of evidence. *International Journal of Environmental Research and Public Health*, 18(23), 12575. https://doi.org/10.3390/ijerph182312575

[16] Dung, P. M. (1995). On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77(2), 321-357. https://doi.org/10.1016/0004-3702(94)00041-X

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.

[17] Chesñevar, C., McGinnis, J., Modgil, S., Rahwan, I., Reed, C., Simari, G., South, M., Vreeswijk, G., & Willmott, S. (2006). Towards an argument interchange format. *The Knowledge Engineering Review*, 21(4), 293-316. https://doi.org/10.1017/S0269888906001044

[18] Agarwal, V., Butler, C., Degenaro, L., Kumar, A., Sailer, A., & Steinder, G. (2022). Compliance-as-code for cybersecurity automation in hybrid cloud. In *2022 IEEE 15th International Conference on Cloud Computing*, 427-437. https://doi.org/10.1109/CLOUD55607.2022.00066

[19] Yanagawa, T., Agarwal, V., Watanabe, Y., Degenaro, L., & Sailer, A. (2024). A secure framework for continuous compliance across heterogeneous policy validation points. In *2024 IEEE 17th International Conference on Cloud Computing*, 176-182. https://doi.org/10.1109/CLOUD62652.2024.00029

[20] Amershi, S., Weld, D., Vorvoreanu, M., Fourney, A., Nushi, B., Collisson, P., Suh, J., Iqbal, S. T., Bennett, P. N., Inkpen, K., Teevan, J., Kikin-Gil, R., & Horvitz, E. (2019). Guidelines for human-AI interaction. In *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 1-13. https://doi.org/10.1145/3290605.3300233

[21] Tabassi, E. (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0)* (NIST AI 100-1). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-1

© Adi Saric. All rights reserved.

This document is protected by copyright. Permission is granted to read, reference, and cite this work with appropriate attribution. No part of this document may be reproduced, adapted, or distributed for commercial use without prior written consent of the authors.