

clara *

RISK.md

A Shared Language for Underwriting

CONCEPT PAPER

VERSION 0.1 / JULY 2026

A proposal for a company-owned, evidence-linked risk record that makes repeated market work reusable without standardizing underwriting judgment.

clarainsure.com/research/risk-md

A Shared Language for Underwriting

A White Paper on Portable, Evidence-Linked Risk Context for Businesses and Insurance Markets

Version 0.1 - Concept Draft · July 6, 2026

Working document for discussion and iteration. Not legal, regulatory, actuarial, claims, or licensed insurance advice.

Executive Summary

Insurance does not need one universal underwriting model. It needs a shared way to describe a business before every company, broker, syndicate, claims team, and agent asks for the same facts again.

RISK.md is one proposal for that shared layer: a company-owned, agent-readable record of exposures, controls, incidents, evidence, unknowns, permissions, and material changes. It would not decide whether a risk is acceptable. It would make the underlying context legible enough for different parties to apply their own judgment without rebuilding the record from scratch.

The proposal combines machine-readable metadata with human-readable context. A useful packet would make clear:

- which legal entity and operating period the record describes;
- which facts are confirmed, estimated, stale, disputed, or unknown;
- which evidence supports each material representation;
- who reviewed or attested to each section;
- what may be shared, with whom, and for which purpose;
- what changed since the last review; and
- which decisions remain with underwriters, claims professionals, regulators, advisors, and accountable humans.

The strategic opportunity is not to standardize underwriting judgment. It is to standardize the risk record that judgment begins from.

1. The Cost of Reconstructing Risk

Business risk context is scattered across applications, spreadsheets, contracts, control reports, loss runs, claims notes, data rooms, security systems, and the memories of people who may not join the next meeting.

Every handoff compresses the story differently. A company restates its operations for a broker. The broker restructures the same facts for a market. Each underwriter asks for overlapping evidence. Claims teams later reconstruct which facts were known, when they changed, and what was represented. Portfolio teams aggregate inconsistent labels after the decisions have already been made.

This fragmentation creates more than administrative cost:

- repeated ingestion delays risk selection;
- stale or unsupported facts survive through copied documents;
- uncertainty is flattened into confident prose;
- different parties cannot tell whether disagreement comes from appetite or inputs;
- sensitive information is overshared because disclosure scope is not explicit; and
- lessons from claims do not reliably return to the underwriting record.

The market should be able to preserve plural judgment without paying to rediscover the same underlying context every time.

1.1 Scope and Definition

For this proposal, a **risk record** is a structured, time-bounded description of an organization, its material exposures, controls, dependencies, incidents, evidence, and known unknowns.

`RISK.md` is the portable human-readable and machine-parseable representation of that record. It is designed to support preparation, review, comparison, and controlled exchange. It is not itself an underwriting decision, a policy, a submission, a warranty, a claims notice, or proof that a control works.

1.2 Relationship to `PASSPORT.md` and `INSURANCE.md`

The three files answer different questions:

| Artifact | Primary object | Question it answers |
|---------------------------|---|---|
| <code>PASSPORT.md</code> | One autonomous agent | Who sponsors this agent, what may it do, and where can its evidence be verified? |
| <code>RISK.md</code> | One organization or defined business system | What risks, controls, dependencies, incidents, evidence, and unknowns describe this organization now? |
| <code>INSURANCE.md</code> | One insurance-readiness context | What insurance-specific facts and evidence are ready for company, advisor, broker, or underwriter review? |

A company `RISK.md` may reference one or many `PASSPORT.md` files as evidence about deployed agents. `INSURANCE.md` may be a purpose-limited module or export from broader risk context. Neither identity evidence nor insurance readiness replaces the full organizational risk record.

2. Why a File Can Change a Workflow

Google Stitch's `DESIGN.md` offers a useful precedent. It pairs machine-readable design tokens with human-readable rationale so intent can travel between people, agents, tools, and sessions. Version control shows what changed. Validation catches structural errors. Exporters translate a stable source into downstream formats.

The important idea is not Markdown. It is durable context with a contract.

Risk information needs the same properties, with stricter evidence and permission boundaries:

- **Company-owned:** portable across advisors and systems rather than trapped in one intake portal.
- **Human and machine readable:** inspectable by executives and professionals while remaining usable by agents and software.
- **Evidence-linked:** material claims point to sources, dates, owners, and limitations.
- **Uncertainty-preserving:** unknowns, estimates, stale facts, and disputes remain visible.
- **Diff-friendly:** material change can be reviewed without rereading the entire company.
- **Permissioned:** the source record can produce purpose-limited views instead of disclosing everything.
- **Extensible:** a thin common core can support class-, jurisdiction-, and workflow-specific modules.

3. The Proposed Record

A first useful `RISK.md` should combine normalized frontmatter with stable Markdown sections.

```

---
schema_version: clara-risk-md-v0
record_id: acme-2026-q3
company_legal_name: Acme Systems Inc.
effective_at: 2026-07-01
risk_context_state: EVIDENCE_LINKED
attestation_status: PARTIALLY_REVIEWED
overall_confidence: MEDIUM
evidence_index: RISK_EVIDENCE_INDEX.md
known_unknowns: "Supplier concentration and open claim reserve require review."
permitted_use: "Internal review and named underwriting pilot only."
---
```

The body should include, at minimum:

- 1 executive risk summary;
- 2 legal entities, ownership, governance, and jurisdictions;
- 3 operations, revenue, products, customers, and material dependencies;
- 4 risk appetite and decision thresholds;
- 5 enterprise risk register;
- 6 controls, policies, tests, and evidence;
- 7 cyber, technology, data, AI, and automation risk;
- 8 contracts, vendors, and third-party obligations;
- 9 financial, capital, and insurance context;
- 10 claims, incidents, disputes, circumstances, and loss history;
- 11 missing information and expected reviewer questions;
- 12 review notes, sharing boundaries, escalation, and attestation; and
- 13 instructions for authorized agents using the record.

3.1 Evidence as the Spine

Each material statement should be traceable to an evidence item, an authorized attestation, or an explicit uncertainty label. An evidence record should identify the source, owner, effective period, sensitivity, facts supported, freshness, and limitations.

The file should not absorb every underlying document. It should point to governed evidence while avoiding credentials, raw secrets, unnecessary personal information, full confidential contracts, and sensitive security details.

3.2 Facts, Analysis, and Decisions

The format should separate three layers:

- **Facts:** what the company and its evidence establish.
- **Analysis:** how a reviewer or model interprets those facts for a defined purpose.
- **Decisions:** appetite, pricing, wording, capital, claims, compliance, or other accountable conclusions.

RISK.md is primarily the first layer. It may preserve attributed analysis, but it should not make one party's conclusion appear to be an objective property of the company.

4. Market Applications

4.1 Shared Ingestion, Vetting, and Clearance

A permissioned, attested module could let a company or coverholder assemble core evidence once, then let each authorized reviewer verify freshness, provenance, and scope. Market participants could add private analysis without forking the underlying facts.

The potential result is less re-keying, faster triage, and clearer disagreement. Two underwriters may reach different conclusions, but they can see whether they began from the same evidence.

4.2 Portfolio Steering and Capital Allocation

Comparable does not mean identical. Stable fields, identifiers, and change histories can make it easier to aggregate exposures, detect drift, trace the source of risk through distribution, and distinguish a portfolio change from a data-quality change.

The record should support portfolio analysis without pretending a common schema creates common appetite.

4.3 Claims and Early-Loss Signals

Structured, dated observations across a claim lifecycle can give early-loss models better inputs and make interventions auditable. Status changes, reserve movements, expert evidence, injuries, litigation milestones, and missing facts can be represented consistently without asking the file to predict severity itself.

The predictive model remains a separate product. RISK.md improves the evidence it is authorized to consume and the audit trail behind its output.

4.4 Company and Advisor Readiness

The same record can help a company prepare for insurance review, lending, M&A, vendor diligence, board reporting, cyber readiness, and AI governance. Purpose-limited exports should disclose the minimum context required for each workflow.

This is why the company must remain the owner of the source context. Portability should increase control, not create a universal risk dossier.

5. What Must Remain Unstandardized

Underwriting discipline depends on plural judgment. Appetite, pricing, wording, capital strategy, fraud detection, sanctions decisions, reserving, claims handling, and regulatory conclusions should not collapse into a public template.

The standard should govern representation, provenance, uncertainty, change, and permission-not compel reliance.

A receiving party must still decide:

- which attestations it accepts;
- which evidence it must verify independently;
- which facts are material to its own appetite;
- what additional questions a class or jurisdiction requires; and
- which decisions require licensed, legal, actuarial, claims, executive, or regulatory authority.

The first implementation should resist a giant ontology. A thin common core with explicit extensions is more adoptable than a schema that attempts to encode every class of business before anyone has used it.

6. Technical and Tooling Direction

The concept is designed around simple, inspectable components:

- **Format:** Markdown with normalized YAML frontmatter and stable section headings.
- **Schema:** machine validation of metadata, evidence references, states, and required sections.
- **Lint:** structural errors, missing sources, stale evidence, conflicting metadata, and unsafe disclosures.
- **Diff:** material changes in exposures, controls, incidents, evidence, permissions, and attestations.
- **Export:** purpose-limited insurance, claims, board, cyber, AI-governance, and diligence views.
- **Audit:** who generated, reviewed, changed, approved, or shared each version.
- **Extensions:** namespaced modules for classes, jurisdictions, and market workflows.

The file is not the database. A hosted system may manage access, evidence, workflows, and audit state while exporting a portable record. The standard should remain usable without requiring that system.

7. Research Plan and Decision Gates

Immediate:

- Test the minimum record against synthetic and permissioned company cases.
- Ask companies, brokers, underwriters, claims specialists, compliance reviewers, and risk leaders which fields they actually reuse.
- Measure whether a shared record improves time to triage, duplicate evidence requests, stale-field detection, reviewer agreement on facts, and handoff quality.
- Map the broader record to the narrower `INSURANCE.md` readiness module and to agent evidence from `PASSPORT.md` where relevant.

Build gate:

- Build broader tooling only after reviewers use the same record in a real, bounded workflow.
- Add fields only when they improve a defined decision or evidence handoff.
- Add automation only where outputs remain attributable, reviewable, and permissioned.

Open-repository gate:

- The current draft is maintained inside Clara Platform, whose active work is narrowly scoped.
- The public white paper and Markdown mirror are the review surface today.
- Create a standalone public specification repository only after external reviewers validate the object, name, and minimum schema.

Stop conditions:

- Existing market standards already provide an adequate portable record.
- Reviewers do not reuse the record across a real workflow.
- Maintenance cost exceeds the duplication it removes.
- The artifact creates false confidence, encourages oversharing, or becomes stale faster than it can be governed.

Success is not the existence of another standard. It is evidence that a shared, company-owned risk record improves a real decision while preserving professional judgment.

8. Project Scope and Maturity

`PASSPORT.md` is a separate open research project with a dedicated public GitHub repository, versioned draft specification, examples, issues, discussions, and a PDF white paper.

`RISK.md` is less mature. It is a Clara Platform concept draft with an internal protocol lineage and a public research paper. It does not yet have a standalone public repository, adopted schema, validator, examples, or market governance process. Those are decision-gated deliverables, not implied capabilities.

This difference is intentional and should remain visible. The publications share an evidence discipline and file-based design philosophy; they do not share the same object, maturity, or adoption claim.

9. Conclusion

Insurance markets do not need every participant to reach the same answer. They need a better way to preserve the facts, evidence, uncertainty, permissions, and changes that each answer depends on.

RISK.md tests one proposition: whether a portable, inspectable record can reduce repeated ingestion, improve handoffs, support better portfolio and claims analysis, and give companies more control over their own risk context.

If the market does not reuse the record, the proposal should stop. If a bounded pilot shows that the same evidence can travel further without losing provenance or context, the next step is to formalize the smallest useful schema and build the tools around it.

The opportunity is not automated underwriting by decree. It is less waste before underwriting begins, better evidence after it does, and a shared memory layer between the businesses that create risk and the markets that evaluate, price, monitor, and respond to it.

References and Further Reading

- [Google Labs: Stitch's DESIGN.md format is now open-source](#)
- [Google Labs: DESIGN.md draft specification and tooling](#)
- [Sønr: Lloyd's Lab Cohort 17 themes](#)
- [PASSPORT.md public research repository](#)
- [Clara: The Economy Will Be Run by Agents. It Must Be Insurable.](#)

RISK.md - A Shared Language for Underwriting

Version 0.1 · July 6, 2026 · Concept draft for market review.

Public Markdown: clarainsure.com/research/risk-md.md