WikiSure Architecture White Paper

Governing Meaning and Determining Which Meaning Applies

Version 1.0 — June 2026
Prepared by SynsureTech

Competent people, working from the same information and the same process, still end up making different decisions — more often than most organizations realize. WikiSure is an architecture for governing meaning. It makes definitions traceable and auditable, and gives people, systems, and AI a working way to tell whether a disagreement is a real conflict or just two valid readings of the same term.

Table of Contents

Executive Summary

Most large organizations have already invested heavily in data quality, governance, enterprise architecture, and now AI. Conflicting decisions still happen anyway. Look closely at where these conflicts actually come from, and a pattern shows up again and again: competent people, correct information, a sound process — and the outcomes still don't line up.

There isn't one cause. Looking at public, authoritative sources across insurance, accounting, privacy law, and financial-crime regulation, two distinct failure modes turn up. Both are real, and organizations typically have no way of telling them apart — let alone resolving either one:

- Meaning Conflict — authoritative sources define the same term differently, and not just from different angles. They are pointing at genuinely different things.
- Applicability Variation — authoritative sources share the same underlying concept, but legitimately apply it under different evaluation frames, for different and equally valid purposes.

Treating the two as one problem is itself risky. Force a single universal definition onto everything, and legitimate Applicability Variation that the organization actually depends on gets destroyed along with it. Wave everything away as "it depends on context," and real Meaning Conflicts get buried — the kind that quietly produce different decisions, audit exposure, and AI errors downstream.

WikiSure is built around that distinction. It governs which meanings are authoritative, classifies how a given divergence should be handled, and resolves which governed meaning applies in a specific context — for people, for systems, and for AI.

None of this claims that universal agreement is achievable, or even desirable. The goal is governed, auditable, context-appropriate resolution — and a working way to tell, case by case, which of the two problems an organization is actually facing.

1. The Problem

1.1 The Hidden Source of Decision Divergence

Two competent teams, same data, same process, same good faith — and they still land in different places. The usual explanations reach for human error, broken process, bad data, or poor communication. Sometimes that's right. Often it isn't the whole story.

There is a failure mode underneath those explanations that tends to get missed: organizations usually have no way to tell whether two authoritative sources mean genuinely different things by the same word, or whether they mean the same thing and are simply applying it through different, equally legitimate, lenses. Until that distinction is made explicit, neither problem gets the response it actually needs.

1.2 The Classification Test

An example only helps if the distinction behind it can actually be tested, not just asserted. So before any examples, here is the test WikiSure applies to any pair of authoritative definitions for the same term:

> *Does the candidate value vary systematically and intentionally by the purpose of the evaluating party, while the underlying entity type stays fixed? Or does the divergence point to a genuinely different kind of thing — a different entity, a different category — rather than a different view of the same thing?*

If the value shifts by purpose while the entity type stays the same, that's an Applicability Variation. Both definitions are correct. What's missing isn't a decision about which one is "right" — it's governance over which frame applies where.

If the divergence points to a different entity type altogether, or the concept simply does not survive translation between the two sources, that is a Meaning Conflict. Only one definition can apply to a given real case at a time, and what is needed is a governance decision about which source controls — not a lookup that adapts itself to context.

This is the test that the Classification stage of the architecture (Section 6.2) turns into a repeatable step. It is stated here, ahead of the examples, because the examples below only make sense once the test that sorts them is visible.

1.3 Meaning Conflict vs. Applicability Variation: Evidence Patterns

Public investigation (Section 2) found both patterns showing up again and again, across domains that have nothing to do with each other. Often a single term contains both at once — which is exactly why classification has to apply to a pair of definitions, not to the term in isolation.

Term	Meaning Conflict pairs found	Applicability Variation pairs found
Coverage	— (insurance-internal investigation; see Section 2.1)	Scope by underwriting context vs. claims context
Risk	ISO 31000 (uncertainty, incl. upside) vs. IRMI (the insured object itself) — different kinds of thing entirely	ISO 31000 vs. NAIC general (uncertainty, loss-only restriction); Solvency II SCR modules (instance

Term	Meaning Conflict pairs found	Applicability Variation pairs found
		vs. capital-aggregation frame)
Customer	ASC 606 (contracting party, can be a legal entity) vs. GDPR (“data subject” — must be a natural person; a corporate customer cannot hold this status at all)	ASC 606 (contract formation) vs. HubSpot lifecycle stage (closed deal) — same entity, same concept, two legitimate gating events

“Customer” shows up in both columns of this table on purpose, not by accident. Classification belongs to a pair of definitions being compared for a specific decision — never to the term by itself. Governing a term correctly means tracking this at the pair level throughout the architecture; see Section 6.2.

2. Public Evidence

Before getting into the architecture, here is the evidence that led to it. Three separate investigations were run against authoritative public sources only — no interviews, no internal documents, and a deliberate attempt to break the underlying thesis rather than confirm it. Between them they cover insurance, cross-industry standards, accounting, privacy law, financial-crime regulation, and commercial CRM practice.

2.1 Coverage Investigation

The first investigation asked a narrow question: could a definition everyone agrees on still produce different, equally legitimate decisions across departments, with nobody actually disputing what the term means?

- Multiple legitimate, simultaneously valid meanings of “Coverage” were found across underwriting, claims, regulatory, and reinsurance contexts.
- Applicability conditions, not definitional disagreement, explained the divergence in decision outcomes.
- Even fully parameterizing the definition — pinning down policy, peril, party, and period until nothing was left open — did not make the divergence go away. What remained unresolved was not a missing fact. It was which evaluation frame (instance, class, portfolio, design space) should govern.

2.2 Risk Investigation

The second investigation went looking for a case that would break the working thesis at the time: that all of this divergence reduces to Applicability. It found one, and the strong version of that thesis did not survive.

- ISO 31000 (“the effect of uncertainty on objectives”), the NAIC Glossary (“uncertainty concerning the possibility of loss by a peril”), and the IRMI Glossary (“the insured or the property to which an insurance policy relates”) were compared directly.
- IRMI’s sense is not a frame variant of the other two — it points at an object, not a property of uncertainty. No amount of parameter binding gets you from “risk = uncertainty” to “risk = the insured car.” That is a genuine Meaning Conflict, not Applicability.
- The same investigation also turned up a clean Applicability case: ISO 31000 and NAIC’s general definition share the same root concept, uncertainty, and differ only in a purpose-driven restriction — loss-only versus including upside.
- None of the sources referenced or disambiguated against any of the others. That is a deeper problem than “no resolution mechanism exists” — none of them even signaled that a resolution was needed in the first place.

2.3 Customer Investigation

The third investigation picked a term with nothing to do with insurance, on purpose, to find out whether the pattern was specific to that industry or something more general.

- ASC 606 (US GAAP): a customer is “a party that has contracted with an entity to obtain goods or services that are an output of the entity’s ordinary activities in exchange for consideration.” A customer can be a legal entity.

- GDPR does not define “customer” at all. Its operative category is “data subject,” and that is defined strictly as a natural person. A corporate customer can never be a data subject; only the individuals connected to it can. That is a categorical Meaning Conflict, not a difference in frame.
- FinCEN’s Customer Due Diligence Rule splits “customer” (the legal entity opening the account) from “beneficial owner” (the natural person who actually owns or controls it) — deliberately, and as a matter of law, because conflating the two is precisely the failure mode the rule exists to prevent. This is not a terminology footnote; it is a regulatory split with real enforcement consequences.
- HubSpot’s lifecycle model defines “Customer” as “a contact or company with at least one closed deal” — the same entity type as ASC 606’s contracting party, just gated by a different and equally legitimate event. Paired with ASC 606, this is a clean Applicability Variation.
- As with Risk: none of these four sources, spanning accounting, EU law, US financial-crime regulation, and commercial software, shows any sign of knowing the others exist.

2.4 What the Evidence Shows Across All Three Investigations

The strong hypothesis this work started with — that decision divergence is never about meaning, only about applicability — did not hold up. It was falsified twice, independently, in domains chosen specifically because they had nothing to do with insurance.

What survived is narrower and harder to attack, and it is what the architecture below is built on:

> *Organizations face both genuine meaning conflicts and legitimate applicability variations, frequently in the same term, and typically have no mechanism to distinguish between them — let alone resolve either.*

That is the mandate behind Sections 5 through 8: not one universal resolution mechanism, but a classification step that decides which of two governance responses a given divergence actually needs, followed by governance and resolution built for each.

3. Why Existing Approaches Fall Short

The tools most organizations already have for managing terminology and metadata are not broken versions of what WikiSure does. They were built to solve different problems, and the evidence in Section 2 sits outside what any of them were ever designed to handle.

3.1 Glossaries

Glossaries store definitions. They have no way of deciding which stored definition fits a given decision, and no way to tell a genuine meaning conflict apart from a legitimate applicability frame sitting right next to it. The NAIC Glossary shows this in practice: “Risk” and “Exposure” appear within a couple of lines of each other, with nothing flagging that the two are not simply interchangeable.

3.2 Data Catalogs

Data catalogs catalog metadata — lineage, ownership, schema — about data assets. Resolving competing meanings of a business term that cuts across several assets, systems, and owners simply is not what they are built for.

3.3 Ontologies and Knowledge Graphs

Ontologies and knowledge graphs represent relationships between concepts with real precision. But representing a relationship and resolving which node in it governs right now, for this decision, are two different operations — and only the first one is what these tools do.

3.4 AI Alone

AI systems generate fluent, plausible-sounding answers. Without governance context behind it, an AI system has no reliable way to know which authorized meaning applies to a given query — it defaults to whatever sense dominated its training data. The Customer investigation suggests that is rarely the sense that actually matters for the decision at hand: accounting, compliance, and privacy obligations look nothing like CRM language.

4. What WikiSure Does Not Claim

A credible architecture is defined as much by its limits as by what it can do. Before getting into what WikiSure does, here is what it deliberately does not attempt — stated plainly, not as something to apologize for.

- WikiSure does not treat every semantic difference as a conflict to resolve. The evidence in Section 2 found legitimate Applicability Variation about as often as genuine Meaning Conflict, and treating one as the other would destroy a kind of coexistence organizations actually depend on.
- WikiSure does not aim for universal standardization across namespaces. Meaning Spaces (Section 7) exist to let multiple authoritative definitions coexist wherever that coexistence is legitimate, rather than pushing every namespace toward one canonical answer.
- WikiSure does not eliminate the need for governance decisions, and never claims to. A Meaning Conflict still needs a person or body with the authority to decide which source

controls. What changes is that the decision becomes explicit, traceable, and applied consistently — instead of being made ad hoc, invisibly, and differently by every team that runs into the term.

5. Semantic Governance

5.1 Definition

Semantic Governance is the discipline of governing meanings — their provenance, ownership, evolution, and applicability. For any given pair of authoritative definitions, it asks one question first: is this a Meaning Conflict, which needs a decision about which source controls, or an Applicability Variation, which needs governance over which frame applies where?

5.2 Objectives

- Govern meanings — establish ownership and authority over which definitions are canonical within a namespace.
- Preserve provenance — record source, authority, and history for every governed definition.
- Detect drift — surface where a previously aligned meaning has begun to diverge across systems or teams.
- Classify divergence — run the test from Section 1.2 before deciding how to govern anything, to find out whether a given divergence is a Meaning Conflict or an Applicability Variation.
- Expose ambiguity — make unresolved or unclassified divergence visible rather than silently defaulting to one sense.
- Enable resolution — return the governed meaning that applies to a specific context, with the reasoning and authority behind it fully traceable.

6. The WikiSure Architecture

6.1 Discovery

Discovery pulls meanings, variants, conflicts, and their provenance out of source documents, systems, and registries.

What comes out the other end is semantic evidence: the raw set of candidate definitions for a term, source and context attached, before anyone has made a judgment about how they relate.

6.2 Classification

This is where the test from Section 1.2 gets applied to every pair of candidate definitions, to work out what kind of difference is actually in front of you. Possible classifications:

- Meaning Conflict — different entity types or categories; at most one definition can govern a given real-world case.
- Applicability Variant — same entity type and root concept, legitimately bound to different evaluation frames for different purposes.
- Alias — different terms, same governed meaning.
- Subtype — a narrower, explicitly bounded specialization of a broader governed meaning (e.g., “Pure Risk” as a subtype of “Risk”).
- Lifecycle Variant — the same entity at different stages of a defined process (e.g., Lead → Opportunity → Customer).

This gets recorded per pair, per namespace comparison, never per term, so a single term like Customer can correctly carry different classifications depending on what it is being compared against — exactly as the evidence in Section 2.3 demands.

The point is to stop organizations, systems, and AI agents from doing either of two things: treating a real Meaning Conflict as if it were harmless Applicability, or treating legitimate Applicability as if it demanded forced standardization.

6.3 Governance

Ownership, authority, provenance, and approval workflows get assigned based on the classification from 6.2. A Meaning Conflict needs a decision about which authoritative source controls in a given namespace, with the rationale and the approving authority on record. An Applicability Variation needs something different: recorded authorization for which frame applies to which purpose, not a ruling on which definition is correct, because in this case, both are.

What you end up with: governed meanings, each with a clear, traceable reason for why it is treated as authoritative in its context.

6.4 Resolution

Resolution determines which governed meaning applies to a specific decision.

Inputs: concept, namespace, context, scope.

Outputs: the resolved meaning, an ambiguity flag where classification is incomplete or contested, a confidence score, and a full provenance trail back to the decision that governs it.

7. Meaning Spaces

7.1 Concept

No organization runs on one universal meaning system. In practice, they operate inside Meaning Spaces — bounded contexts like insurance, compliance, claims, underwriting, or finance — and each one can legitimately hold its own authoritative definition for a term the others also use.

Meaning Spaces are what make legitimate coexistence possible without forcing standardization. An Applicability Variation between two Meaning Spaces is not an error to clean up — it is a relationship that needs governing and a clear paper trail, not elimination.

8. Namespace-Scoped Resolution

8.1 Why Namespace Matters

The same term — coverage, claim, risk, customer — can exist across several authoritative namespaces at once. Namespace-scoped resolution is designed to retrieve, deterministically, whichever definition is canonical within a given namespace.

8.2 Resolver Architecture

Inputs: identifier, namespace.

Outputs: resolved meaning, provenance, confidence, governance status.

8.3 Current Capability and Roadmap

As it stands today, namespace-scoped resolution can tell you which definition is canonical within a given namespace. It cannot yet tell you, on its own, whether the divergence between two namespaces is a Meaning Conflict or an Applicability Variation — that still requires the Classification step from Section 6.2 to be applied and recorded separately.

Closing that gap is a near-term priority: getting the resolver itself to tell “a different governed meaning exists elsewhere” apart from “a different evaluation frame legitimately applies elsewhere.” That means an explicit frame or evaluation-unit parameter, separate from namespace, that the resolver can query directly instead of relying on classification metadata sitting outside the resolution path.

9. Implications for AI

AI systems inherit whatever semantic ambiguity sits in the sources they are trained or grounded on. Without governed meaning behind them:

- Retrieval fails — the wrong sense of a term is returned with no signal that alternatives existed.
- Reasoning drifts — an agent silently carries one sense of a term across a task where a different sense was required at a later step.
- Decisions diverge — two AI-assisted workflows reach different, individually defensible conclusions for reasons neither workflow can surface.
- Auditability collapses — there is no trail showing which meaning was used, why, or under what authority.

Semantic Governance is the missing control layer here — not by handing AI one universal definition to fall back on, but by making the classification and resolution machinery in Sections 6 and 8 something an AI system can actually query at decision time.

10. Implications for Enterprise Architecture

Semantic ambiguity behaves a lot like technical debt. It is invisible right up until a decision depends on it, and the cost of resolving it grows with every system and process built on top of it in the meantime.

WikiSure introduces three architectural capabilities to address this directly:

- Semantic Governance — ownership, provenance, and authority over canonical definitions.
- Meaning Resolution — deterministic retrieval of the governed meaning for a given namespace and context.
- Applicability Determination — the explicit classification and governance of legitimate frame variation, distinct from and complementary to Meaning Resolution.

11. Conclusion

Organizations do not fail simply because meanings differ. Difference on its own is normal — and where it is a legitimate Applicability Variation, often the right outcome.

They fail when there is no way to tell a genuine Meaning Conflict apart from a legitimate Applicability Variation, no governance fitted to whichever one is actually in play, and no way to determine which governed meaning applies to a specific decision — for humans, for systems, and increasingly, for AI.

WikiSure provides that mechanism. The goal was never universal agreement. It is governed, auditable, context-appropriate meaning resolution — and an explicit classification step that decides, case by case, which kind of governance a given divergence actually needs.

Related Materials

WikiSure Use Case Map (Sales / Investor Appendix) — extends the evidence in Section 2 into a tiered map of additional use cases (evidenced, plausible, vision), without widening the claims made in this paper.

MeaningRoom Applicability Matrix series (Category Evidence Artifacts) — independent, illustrative evidence pieces published under MeaningRoom, demonstrating that the Meaning Conflict / Applicability Variation distinction recurs outside this paper's three core investigations. These are category evidence, not claims about WikiSure as a mechanism.