Enterprise Reasoning for Decision Intelligence

An open standard for interoperable, explainable, reusable reasoning

September 2026

Authors: Bob Kruger, Paul Stanton, Ramesh Parameswaran, Greg Coquillo, Sube Chatterjee, Vidya Subramanian, Dan Wright

Executive Summary

AI models respond instantaneously to questions and are increasing the speed and volume at which organizations generate analyses and answers. Yet the output of large language models is designed for conversation and for AI models to consume, but not for enterprise decision-making. Verbose unstructured text is difficult for software to digest, hard for teams to collaborate on, and inadequate for decision-making. Humans are being trained by AI models to auto-accept their verbose explanations leading to abdication of human reasoning.

As AI does more of the reasoning behind business decisions, the people who understand the business become less involved—eroding trust in the decisions that result.

What is needed is reasoning that is initially proposed by AI models, elaborated by humans, explainable to teams, persisted for reuse, and available for review and approvals. As a result, reasoning becomes a durable enterprise asset, reusable across models and teams. Competitive advantage goes beyond access to a particular model and instead is created from the organizational learning loop built on structured reasoning.

This paper proposes **Enterprise Reasoning (ER)** defined as a set of capabilities and an open, vendor-neutral, interoperable standard for representing reasoning as a structured artifact. Rather than treating reasoning as disposable chat, Enterprise Reasoning treats it as a reusable, explainable, structured versioned asset that is shared among people, software tools, and AI models.

The goal of ER is to enable enterprises to explain not only **what** decisions were made, but **why** they were made, **who** approved them, **what** they cost, and **which** outcomes were influenced.

1. The Problem

AI models can reason, but enterprises currently have no common way to capture, understand, and build upon that reasoning. Organizations also have no way to own the IP behind the reasoning they created and in fact relinquish ownership to the AI models.

Today, AI reasoning is expressed primarily as verbose text, code, skills, projects, or ephemeral interactions. None of these provide a common, durable representation of reasoning that both humans and software can readily inspect, modify, govern, and reuse. Text is accessible to people but difficult for software to structure and build upon. Code can be executed and reused by software, but is increasingly difficult for humans, particularly non-technical business users, to understand and modify. As autonomous agents become more prevalent, reasoning that might previously have been expressed in documents is increasingly embedded in application and agent code, and chat.

This creates several problems for enterprises.

First, AI model **reasoning is difficult for humans to interrogate and participate in.** Domain experts may have the knowledge and judgment required to improve an AI-generated decision, but today's interfaces often reduce their role to reviewing a final answer or approving an action. This makes it difficult for experts across domains to examine the evidence and assumptions, contribute their judgment, and collaboratively improve the reasoning before a decision is made.

Second, reasoning is not interoperable. Without a common representation, reasoning becomes coupled to the application, agent framework, or AI model that produced it. Each application develops its own representation, limiting the ability of humans and software to move reasoning between systems or use different AI models and tools. This creates a form of vendor and application lock-in at the reasoning level.

Third, reasoning is difficult to govern. Effective governance requires more than approving an AI-generated result. Enterprises need to know how a result was produced, what data and assumptions were used, what decisions were made by humans, and how the reasoning has changed over time. This requires reasoning to be explainable, persistent, attributable, and capable of being versioned and reviewed.

Fourthly, with the advent of autonomous agents, the latest expertise and context in an organization may not be used by AI to advance reasoning.

Finally, enterprises are accumulating AI interactions rather than organizational intelligence. An AI conversation may contain valuable business reasoning, but that reasoning generally remains trapped in the conversation or application that produced it. It cannot easily be preserved as an enterprise asset, improved collaboratively, or reused in future decisions. The intellectual capital generated through AI therefore fails to compound.

All of these problems ultimately lead to lack of trust in AI outputs.

Recent research highlights the consequences of this gap. A June 2026 study by researchers from Cambridge University and China Agricultural University conducted 280 research runs comparing an unconstrained multi-agent system with the same system incorporating three human decision gates (See https://arxiv.org/pdf/2606.12848). The unconstrained system experienced critical failures in 72% of runs, compared with 16% when humans participated at decision gates. The study provides experimental evidence that human judgment can be more effective when incorporated into the reasoning process rather than applied only as validation of the final result.

HFS research published a study titled *Humans at the Helm of AI* in April 2026 (https://www.hfsresearch.com/research/humans-at-the-helm-of-ai/) where they surveyed 500 executives across Global 2000 organizations and they concluded that "Fifty-three percent (53%) name human-in-the-loop as their primary governance mechanism. Only 18% can actually interrogate the reasoning behind what they are approving. The thing enterprises trust most to keep them safe is the thing they have invested in least." The wording "... only 18% can interrogate …" is significant since that indicates an inability to interrogate the reasoning as opposed to the unwillingness to do so. The issue is therefore not simply whether enterprises want humans involved; many already do. The problem is that they lack the capability to meaningfully interrogate and participate in AI reasoning.

A further study by researchers at Saarland University in Germany found that autonomous AI systems reduce the effort people devote to monitoring those systems and shift responsibility away from themselves (See (https://www.sciencedirect.com/science/article/pii/S2949882126000885?utm_source=chatgpt.com). This reinforces the need for humans to remain meaningful participants in reasoning and decision making rather than becoming passive approvers of AI-generated outcomes.

Together, these findings point to a fundamental gap. Enterprises need humans and AI to reason collaboratively, but the systems through which they interact do not provide a common representation of that reasoning. Without such a representation, human judgment remains difficult to incorporate, reasoning cannot readily move between systems, governance is difficult, and valuable organizational knowledge disappears after each AI interaction.

As Microsoft CEO Satya Nadella observed, "The real opportunity is not in picking the best model but instead in building a learning loop on top of models where human capital and token capital compound."

For that learning loop to operate at an organizational level, reasoning must become an enterprise asset: something that can be captured, understood, interrogated, versioned, improved, governed, and reused across people, applications, and AI models.

Enterprise Reasoning proposes a common representation for that purpose.

2. What is Enterprise Reasoning

Enterprise Reasoning proposes that reasoning become a **first-class enterprise artifact**, much like source code, documents, or datasets. Enterprise Reasoning standardizes the representation of reasoning so that people, software tools, and AI models can exchange, collaborate on, govern, validate, and build upon it.

A standardized representation of reasoning helps remove the obstacle for humans to "interrogate reasoning" described in the HFS study above. The standard format also enables software tools to collect the increasing volume of AI-created reasoning and dashboard it for humans to interrogate. Further, this standardization supports workflows involving "human decision gates" in the middle of AI reasoning as described in the research experiment above. Software tools can easily create workflows and dashboards requiring such human decision gates as long as they can parse reasoning steps.

Rather than storing chats, code, and the final answer, an Enterprise Reasoning artifact includes a structured format described in the structured reasoning section. The proposed standard of representation of reasoning enables six capabilities. It does not standardize how AI models reason or how software implements these capabilities. The six capabilities are:

Structured Reasoning

Human AI collaborative reasoning

Reasoning Lifecycle Management

Reasoning Across Multiple Information Sources

Progressively Evolving Semantic Intelligence

Attribution of reasoning to decisions and outcomes

3. Structured Reasoning

Structured Reasoning is a representation of reasoning using a defined information model rather than free-form text or code. The information model defines the concepts, relationships, and properties needed to describe how a problem is being addressed and how a conclusion or result is reached.

The information model represents the context of the problem, the information and evidence available, the assumptions being made, the steps performed, the decisions made, and the resulting outputs. Each of these is represented using standardized terms and relationships so that AI models, humans and software can interpret the reasoning consistently.

For example, reasoning over structured enterprise databases may represent the databases, tables, columns, and other data objects used as sources; the semantic meaning of the information; assumptions and constraints; individual reasoning steps like joins, and filters; calculations or transformations; evidence supporting each step; and the resulting outputs.

A reasoning step may be produced by an AI model, a human, or a software tool. A structured representation can therefore capture collaborative reasoning rather than treating AI-generated reasoning as fundamentally different from human reasoning.

The Enterprise Reasoning Information Model could include, for example:

1. **Problem context**
    - The question, problem, objective, or decision being addressed
    - Relevant business context and semantic meaning
2. **Sources**
    - Sources of data or information
    - References to databases, tables, documents, files, or other information objects
    - The specific information used by each reasoning step
3. **Assumptions and constraints**
    - Assumptions made during reasoning
    - Business rules, constraints, and other conditions that affect the reasoning
4. **Reasoning steps**
   Each step could contain:
    - **Step ID -** Unique identifier so other steps can reference it cleanly in
    - **Timestamp** — when it ran (useful for latency debugging and for time-sensitive data)
    - **Status** — success / failed / retried / superseded
    - **Producer** — AI model, human, or software tool
    - **Title** — description of the step
    - **Type** — such as calculation, data transformation, code execution, inference, human judgment, or decision
    - **Details** — the logic or operation performed
    - **Evidence or argument** — why the step is justified
    - **References** — the data or information used
5. **Policies and governance**
    - Information policy compliance
    - Data access and authorization
    - Privacy controls, such as whether PII has been hidden or transformed
    - Human approvals and review status
6. **Results and outputs**
    - Where results are stored or presented
    - References to outputs produced by the reasoning
    - Relationships between intermediate and final results
7. **Reproducibility**
    - Information required to execute the reasoning again
    - The tools, models, code, queries, and parameters required
    - The conditions under which the same reasoning should produce equivalent results from the same source information
8. **Version and lifecycle**
    - Version of the reasoning

- Changes between versions
- Approval stage and history
- Provenance of changes

This information model establishes reasoning as a first-class enterprise artifact rather than an ephemeral response from an AI system. It allows reasoning to be read by a person, processed by software, modified by another AI model, reviewed by a subject matter expert, or reused as the starting point for another decision.

The standard should define the information model independently of its serialization format. JSON, YAML, or other syntaxes could be used to represent the same underlying reasoning. Multiple syntaxes may be useful because different AI models, software systems, or human workflows may perform better with different representations. Where multiple syntaxes are supported, they should be deterministically convertible so that the underlying reasoning remains portable. The information model should be model, agent, and implementation agnostic, enabling reasoning artifacts to remain reusable as the underlying AI models, agent frameworks, tools, and enterprise systems evolve or are replaced.

The standard may also define conventions for naming concepts and properties. Naming is important both for human collaboration and for AI systems that must interpret and modify existing reasoning. As AI models evolve, different representations or naming conventions may prove more effective for particular use cases. A standardized information model provides the stable semantic foundation while allowing its representation to evolve.

The Enterprise Reasoning Information Model should ultimately support domain-specific extensions. The core model can describe common concepts across reasoning activities, while extensions can address particular domains and use cases, for example, reasoning over structured enterprise data, reasoning over unstructured documents, healthcare, banking, or financial planning and analysis.

In this way, Structured Reasoning provides a common language through which humans, AI models, and software can read, create, interrogate, modify, and reuse reasoning.

4. Human-AI Collaborative Reasoning and Agent-Agent Collaborative reasoning

Agentic workflows increasingly enable AI systems to plan and execute multi-step tasks autonomously. However, an agentic workflow often treats the reasoning underlying those tasks as something that is generated and executed by the AI. In many cases, there is an underlying sequence of decisions, assumptions, calculations, and actions that could instead be captured as an explicit, reviewable reasoning artifact. Once that reasoning has been approved, an agent can execute it rather than independently recreating the reasoning each time.

Today, a typical AI workflow places humans primarily at the end of the process, where they validate an AI-generated recommendation or approve an action. Enterprise Reasoning proposes a different model: humans and AI collaborate throughout the reasoning process itself. Domain experts may modify assumptions, introduce new evidence, override a recommendation, interpret unusual business conditions, or apply regulatory and organizational knowledge while reasoning is taking place. To enable this collaboration, reasoning needs to be represented as a sequence of discrete, understandable steps. Humans and AI should be able to:

- Introduce a new reasoning step between existing steps
- Modify or delete an existing step
- Understand the purpose, inputs, evidence, and outcome of each step
- Review and approve changes to reasoning
- Preserve the resulting reasoning as a versioned enterprise artifact

This creates an important distinction between agentic execution and collaborative reasoning. An agent can execute an approved reasoning plan, while humans and AI can collaboratively develop, modify, and improve that plan. The reasoning therefore becomes persistent and reusable rather than being regenerated as an opaque part of every agent interaction.

Enterprise Reasoning brings humans into the reasoning process itself, allowing human judgment and AI generated reasoning to compound rather than treating humans solely as validators of AI output.

5. Reasoning Lifecycle Management

Reasoning should be managed as a first-class enterprise artifact, much like software code. It should be possible to create, test, review, approve, version, publish, execute, modify, and reuse reasoning rather than regenerate it from scratch each time a related question is asked.

Software engineering and DevOps have established lifecycle practices for managing artifacts that need to evolve reliably. Enterprise Reasoning applies similar principles to reasoning:

- **Versioning** — maintain identifiable versions of reasoning
- **Staging and approval** — move reasoning through states such as draft, testing, review, and approval
- **Diffs** — identify and understand changes between versions
- **Rollback** — return to a previously approved version
- **Repository** — store, discover, create, modify, and retire reasoning artifacts
- **Execution** — run the defined reasoning against its source information
- **Provenance** — maintain the history of how reasoning was created and modified
- **Triggers** - changes to reasoning based on changes to schemas, metrics, etc.

Lifecycle management allows an organization to treat approved reasoning as reusable organizational intelligence. A new analysis should not require an AI model to rediscover reasoning from the potentially massive context used to produce the original reasoning. An approved reasoning artifact can instead serve as structured context for the new analysis, allowing humans or AI to modify only the reasoning that needs to change. This creates several benefits:

- **Continuous improvement** — organizations can refine reasoning over time and deploy improved versions
- **Reuse** — new analyses can begin from previously approved reasoning
- **Consistency** — related problems can use harmonized reasoning across teams
- **Reproducibility** — organizations understand and rerun reasoning with the same results
- **Operationalization** — approved reasoning can become repeatable workflows (software or agent driven workflows) rather than one-time AI interactions
- **Agentic execution** — autonomous agents can execute workflows composed of multiple approved reasoning artifacts
- **Reduced token consumption**— reusing established reasoning can reduce the amount of AI context and tokens required for each new analysis

The result is a shift from AI generating reasoning on demand to organizations accumulating and improving reasoning over time. Previously validated reasoning becomes context for future problems, allowing human expertise and AI generated reasoning to compound rather than repeatedly starting from scratch.

Reasoning lifecycle management therefore provides the operational foundation for Enterprise Reasoning: reasoning can be developed collaboratively, validated, deployed, reused, measured, and continuously improved as an enterprise asset.

6. Reasoning Across Multiple Information Sources

AI is capable of reasoning across multiple sources of enterprise information—databases, documents, and other systems—without requiring organizations to first centralize or transform all information into an “AI-ready” data environment. Data can remain within its existing governed boundaries while AI reasons over the information relevant to particular questions.

In structured databases, Enterprise Reasoning combines human-specified semantics with AI inference. Database schemas, table and column names, primary and foreign keys, and other database structures already capture significant human knowledge about the meaning and relationships of data within a database. AI can use this knowledge directly and can increasingly infer relationships across independently designed databases where explicit relationships do not exist.

This capability extends reasoning beyond the information and relationships that have previously been modeled into data warehouses and pipelines. Data pipelines necessarily reflect decisions and assumptions made when they were created, such as what information was important, how it should be transformed, and how sources should be joined. Those decisions remain valuable, but they cannot anticipate every future business question. AI can reason over additional source information in response to a new question, identifying relevant data, joins, and relationships that may not have been captured previously.

Enterprise Reasoning therefore does not replace or preclude warehouses, pipelines, or other curated data environments. It extends reasoning across the broader enterprise information landscape, allowing the resulting reasoning to be captured for human review, reuse, and organizational intelligence.

AI-ready data is not a prerequisite for AI reasoning. Enterprise Reasoning allows AI to reason over existing enterprise information, within its governed boundaries, using both the human intelligence already embedded in the data and AI's ability to infer relationships.

7. Progressively Evolving Semantic Intelligence

Enterprise semantics should not be treated as something that must be completely defined before AI can reason. Organizations already possess substantial semantic intelligence developed over years by business and technical teams. It is embedded in metrics, reports, database schemas, business glossaries, applications, business rules, recently developed ontologies, knowledge graphs, and the knowledge of domain experts.

Enterprise Reasoning extends this existing semantic intelligence rather than replacing it. AI reasoning across enterprise information can identify new relationships, metrics, categories, classifications, and business concepts that build upon what the organization already knows. Human experts can evaluate these discoveries, refine their meaning, and approve them for use. Once approved, the resulting semantic definitions become reusable organizational intelligence and provide context for subsequent reasoning.

Semantics can therefore be discovered, validated, harmonized, and progressively improved through use and enhancing existing semantic work.

Enterprise Reasoning treats semantic definitions as managed, versioned assets that evolve through this process. Examples include:

- Metrics and measures — such as `gross_margin_pct`
- Business concepts and categories — such as `strategic_account` or `high_value_customer`
- Business rules and classifications
- Column names, table names and their relationship with metrics
- Code, formulas, and other executable definitions

- Relationships between semantic concepts

A semantic definition can have its own lifecycle, for example:

```
gross_margin_pct v1.3 – Approved
gross_margin_pct v1.4 – Tested
```

This lifecycle enables organizations to:

- Extend existing semantics as new information and business requirements emerge
- Progressively improve semantics through actual use and human-AI collaboration
- Harmonize semantics across business units, applications, and AI systems
- Prevent semantic drift by ensuring that an approved definition has a consistent meaning when reused
- Preserve provenance by recording who or what proposed, modified, and approved a definition
- Associate reasoning with semantics by recording which definitions and versions were used to produce a result
- Prioritize approved definitions when AI models encounter competing or ambiguous interpretations

This creates a feedback loop between reasoning and semantic intelligence. AI can discover potential meaning from enterprise information; humans can validate and improve it; and approved definitions can then guide future AI reasoning.

Over time, the enterprise therefore accumulates not only reusable reasoning, but an evolving body of semantic intelligence that builds on the organization's existing knowledge and becomes increasingly consistent and valuable with use.

Enterprise Reasoning turns semantics from a static prerequisite for AI into a living, governed enterprise asset that evolves through human-AI collaboration.

8. Attribution of Reasoning to Outcomes

AI-generated decisions create a fundamental accountability challenge: organizations need to know not only what decision was made, but how the decision was reasoned, who contributed to it, and what happened as a result.

Research from Saarland University indicates that autonomous AI systems can impact accountability negatively by reducing the effort people devote to monitoring those systems and shifting responsibility away from themselves. This reinforces the importance of keeping humans meaningfully involved in reasoning and maintaining a clear connection between reasoning, decisions, and outcomes.

Enterprise Reasoning creates this connection by persisting the reasoning behind decisions and attributing contributions throughout the reasoning process. Each reasoning step can identify its producer—AI model, human, or software—along with the evidence used, assumptions made, approvals received, and changes introduced.

The resulting reasoning artifact can then be associated with the decision and, ultimately, with the business outcome.

This enables organizations to answer questions such as:

- What reasoning led to this decision?
- Which AI models, people, and systems contributed?
- What human judgment changed or approved the reasoning?
- What evidence supported the decision?
- Which version of the reasoning was used?
- What did the reasoning cost to produce?
- What business outcome resulted?

Attribution improves both trust and accountability. Trust is strengthened when human expertise and judgment are visible alongside AI-generated reasoning. Accountability is strengthened when decisions and outcomes can be traced back to the reasoning and contributions that produced them.

More importantly, attribution creates a feedback loop for improving decision intelligence. Organizations can compare reasoning and decisions with their actual outcomes, identify which reasoning was effective, and improve future reasoning accordingly.

Enterprise Reasoning therefore connects AI reasoning to human judgment, decisions, and measurable business outcomes, allowing organizations to learn systematically from the decisions they make rather than simply accumulating AI-generated answers.

9. An Example - “Which wholesalers to cut”

An example question for a manufacturer of a popular product to ask every quarter could be “Which wholesalers should we prioritize or cut?”. The data, customers, and rules for keeping wholesalers may change every quarter. The first time a user asks the model the question, a reasoning plan is created by the model. Every quarter, when the question is asked, the same plan is reused to create the next version based on changes in requirements. The reasoning could be structured as below:
- Decision: "Which wholesalers to cut this quarter"
- Source databases: (Oracle Financials, Postgres reseller database)
- Step 1 (AI):Join revenue, cost tables from Oracle to Postgres wholesalers table

- Step 2 (AI): Compute margin for 1/1/26 to 3/31/26 grouped by wholesaler id for each wholesaler
- Step 3 (AI) Filter to get wholesalers below the 12% gross-margin floor
- Step 4 (Human): exclude 6 in hurricane-affected territories — disruption is temporary
- Step 5 (AI): 14 of the remaining 21 are below floor on product mix alone, renegotiate, don't cut
- Definitions used: Gross margin version 2.1.0
- Results table: In SQL Server analytics DB

This plan could be expressed in JSON form as follows, which in turn could be parsed by a tool to make it human-friendly :


```
{
  "er_version": "0.1",
  "plan": {
    "id": "wholesaler-review",
    "decision": "Which wholesalers should we prioritize or cut?",
    "version": "4.2.0",
    "previous_version": "4.1.0",
    "stage": "approved",
    "period": "2026-Q2"
  },

  "sources": [
    "Oracle Financials",
    "Postgres Reseller Database"
  ],

  "semantics": {
    "gross_margin_pct": {
      "definition": "(net_revenue - landed_cost) / net_revenue",
      "version": "2.1.0",
      "status": "approved",
      "owner": "finance"
    }
  },

  "reasoning": [
    {
      "step": 1,
      "produced_by": "AI",
      "kind": "join",
      "details": "Join revenue and cost data from Oracle Financials with the Postgres wholesalers
table."
    },
    {
```

```
      "step": 2,
      "produced_by": "AI",
      "kind": "calculation",
      "details": "Compute gross margin for 2026-04-01 through 2026-06-30 for each wholesaler."
    },
    {
      "step": 3,
      "produced_by": "AI",
      "kind": "filter",
      "details": "Identify wholesalers below the 12% gross-margin floor.",
      "result": "27 wholesalers"
    },
    {
      "step": 4,
      "produced_by": "Human",
      "kind": "judgment",
      "details": "Exclude 6 wholesalers in hurricane-affected territories.",
      "rationale": "The disruption is temporary and should not be treated as a performance signal."
    },
    {
      "step": 5,
      "produced_by": "AI",
      "kind": "analysis",
      "details": "Of the remaining 21 wholesalers, 14 are below the margin floor primarily because
of product mix.",
      "recommendation": "Renegotiate rather than terminate."
    }
  ],

  "result": {
    "terminate": 7,
    "renegotiate": 14,
    "results_location": "SQL Server Analytics Database"
  },

  "attribution": {
    "human_reviewers": ["sales leadership"],
    "approval_status": "approved"
  },

  "changes_from_previous_version": [
    "Updated analysis period to 2026-Q2",
    "Added hurricane-affected territory exclusion",
    "Applied gross_margin_pct semantic definition version 2.1.0"
```

```
  ]
}
```

10. Why an Open Standard?

Open standards are emerging across the enterprise AI ecosystem. The Model Context Protocol (MCP) standardizes how AI models access tools and information. Agent-to-Agent (A2A) protocols standardize how autonomous agents communicate with one another.

Enterprise Reasoning addresses a different layer: the reasoning itself.

A common information model for reasoning enables humans, AI models, and software to create, understand, modify, and reuse the same reasoning artifacts. This makes reasoning a shared enterprise capability rather than a proprietary feature of an individual AI platform.

An open standard for reasoning enables:

- Human-AI collaboration — people can participate directly in structured reasoning
- Broader community participation — non-coding professionals can understand and contribute to reasoning without having to work through application code
- Interoperability — reasoning can move between AI models, applications, and software platforms
- Reuse and organizational intelligence — approved reasoning can become a shared enterprise asset
- An interoperable software ecosystem — vendors can build tools that create, consume, inspect, govern, and execute reasoning artifacts
- Enterprise decision intelligence — organizations can continuously develop, reuse, and improve the reasoning behind their decisions

A data format such as JSON provides syntax, but syntax alone does not provide interoperability. Systems must share an understanding of what the elements of a reasoning artifact mean. An open standard therefore needs to define the **information model and semantics of reasoning**, while allowing that model to be represented in different syntaxes and implemented by different systems.

An open Enterprise Reasoning standard would allow the ecosystem to build around a common representation of reasoning rather than creating isolated proprietary implementations. It provides the foundation for reasoning to become a portable, collaborative, and reusable enterprise asset—and ultimately a capability through which organizations can develop and improve decision intelligence.

To learn more, visit enterprisereasoning.org or email contact@enterprisereasoning.org.