

Beyond Data: Why Knowledge Representation Is an Enterprise Architecture Concern

Building Interoperability on Purposeful, Governed, and Plural Representations

Enterprise Knowledge Representation — Article 1

Jean Vieille

23 September 2026

DOI: 10.5281/zenodo.22918368

Abstract

Operational data are increasingly accessible through APIs, Unified Namespace architectures, metadata and catalogue platforms, information models, data platforms, ontologies, and knowledge graphs. Yet accessibility alone does not ensure reliable interpretation and use across enterprise boundaries. Effective interaction also depends on operational knowledge about identity, relationships, authority, provenance, validity, capabilities, constraints, and purpose. This article introduces Enterprise Knowledge Representation (EKR) as an architectural capability for creating, relating, and governing purposeful, plural representations of such knowledge. EKR builds on existing authoritative models rather than replacing them and creates representation readiness for interoperability, semantic mediation, and composable participation. The article also introduces the emerging EKR Architectural Guidance and its open development process.

Introduction

An enterprise takes shape, operates, and evolves through interactions. People coordinate or execute activities and make decisions; equipment transforms materials and energy; applications exchange information and services; organisations establish commitments; and increasingly software and AI agents interpret context and initiate or support actions. These interactions are what turn resources, data, and expertise into operational outcomes.

Data drives or supports many of those interactions, but data alone is rarely enough once participants no longer share the same local assumptions. For many practitioners, part of this challenge first appears as a metadata problem: definitions, ownership, lineage, classifications, quality, provenance, and policy provide essential context around data. Metadata is an important part of the picture, but not a synonym for it. Operational interaction can also depend on identities, distinctions, relationships, roles and capabilities, responsibilities, constraints, authority, validity, and purpose. Representation matters because it makes selected knowledge communicable and usable for interaction, rather than merely adding description around data. Here, a representation is the explicit form in which selected operational knowledge about a defined scope is made available for a particular purpose — what is being referred to, how relevant things relate, what meaning and authority apply, and under what conditions the knowledge is valid and usable.

Enterprise systems have become increasingly effective at making operational data and context available. Values can be published through APIs, brokers, and Unified Namespace patterns; information models can add structure and semantics; metadata and catalogue platforms can manage definitions, ownership, lineage, classifications, and governance context; data platforms can contextualise and persist information; ontologies and knowledge graphs can formalise concepts and relationships. These are important advances and form part of the solution.

The difficulty becomes particularly acute when information and knowledge must support interactions across systems, sites, disciplines, organisations, or lifecycle stages. Availability alone does not ensure that a distant participant can interpret a representation correctly, understand its authority and validity,

relate it to another representation, or determine what actions and interactions it can safely and usefully support.

The practical objective is to make selected operational knowledge sufficiently explicit, related, and governed so that people, applications, equipment, organisations, and agents have a reliable basis for effective and adaptable interactions. This does not require one universal enterprise model. Better governed representations can support discovery and onboarding, make compatibility assessable, and create readiness for semantic mediation and composable participation.

The Enterprise Knowledge Representation — Working Definition V0.1 defines EKR as follows:

“Enterprise Knowledge Representation, or EKR, is the enterprise capability to create, relate, govern and evolve purposeful representations of selected operational knowledge, so that people, applications, equipment, organisations and agents have a structured and governed foundation from which semantic and operational interoperability can be progressively established across systems, sites, disciplines, organisations and lifecycle stages.” [1]

Two terms are sufficient for the discussion that follows. An EKR is a particular explicit, partial, and purpose-specific representation. The EKR capability is the broader enterprise capability that creates, relates, governs, validates, evaluates, exposes, and evolves multiple such representations. [1] An architecture does not constitute an EKR capability merely because it contains metadata, information models, ontologies, or knowledge graphs. For a declared purpose and scope, a defensible EKR claim should be able to show which representations are needed, how their identities and relationships are governed, where authority and provenance reside, how validity and evolution are handled, and how those representations support the intended interactions.

At this stage, EKR is a proposed enterprise-capability framing. Its architectural guidance is still under development. The Working Definition provides the conceptual anchor; the public EKR workspace is developing architectural questions, evidence, patterns, and implementation experience into progressively stronger guidance.

This article therefore has two purposes. First, it explains why governed enterprise representation deserves explicit architectural treatment, using cross-boundary interaction as the stress case in which the need for those requirements becomes most apparent. Second, it connects that framing to the open development process now in place, so that the distinctions proposed here can be challenged, tested, and refined. The intended value is practical: better interoperability, discovery and onboarding, compatibility assessment, readiness for semantic mediation, and more composable participation — without turning those downstream capabilities into intrinsic EKR functions.

Although the EKR framing is intended to apply beyond industrial settings, it has been developed primarily from the perspective of complex industrial and automation environments. These environments provide demanding stress cases because they combine long asset lifecycles, multiple engineering disciplines, heterogeneous information models, organisational handovers, and strong IT/OT boundaries. The broader applicability of the framing should therefore be tested in other enterprise contexts rather than assumed.

AI does not create the representation problem, but delegation to software and AI agents raises its stakes. Human specialists can sometimes recover missing context through experience, discussion, or investigation; an agent may instead produce a plausible interpretation without knowing that an identity, authority, or provenance assumption is wrong. Before interpretation or action is delegated, the operational context needed for that decision should therefore be explicit, machine-interpretable, and governed. Representations created or amended by AI agents should retain explicit provenance, authority, and validation status rather than silently acquire operational authority. EKR does not make an agent trustworthy by itself; it is intended to make the knowledge on which such delegation depends more explicit and governable.

1. The problem becomes harder at the boundary

Consider a simple operational value:

Transporting that number is easy. The originating system may also expose rich contextual information: the object concerned, unit, timestamp, quality, source, engineering structure, process role, or a domain model that gives those elements a precise local meaning.

Within a well-bounded domain, explicit modelling, information models, and ontologies can provide a shared conceptual frame for the people and systems that adopt them. A local model may itself qualify as an EKR if it satisfies the EKR definition and the requirements applicable to its declared purpose and scope; otherwise it can remain an authoritative source representation from which the EKR capability derives selected knowledge without transferring its local authority.

The difficulty appears when operational knowledge must cross a boundary between domains whose representation structures, vocabularies, granularities, authorities, or purposes are not fully shared.

For example, one domain may expose an OPC UA information model while another relies on RDF/OWL representations. Both may be internally coherent and useful. The problem is not that either representation lacks semantics, nor that the EKR framing requires every term to be translated into one enterprise vocabulary. The problem is that the enterprise needs a governed way to make those representations identifiable, relatable, and usable together while preserving their respective authorities and differences. [4][7]

Why this matters is operational. When identities, authorities, representation differences, capabilities, and interaction conditions are explicit and governed, downstream mechanisms can more reliably establish semantic correspondences, assess compatibility, support discovery and onboarding, assist people or AI in interpretation, and prepare participants for composable interaction. The EKR capability is intended to create that representation readiness: semantic mediation, compatibility decisions, coordination, and execution remain downstream capabilities.

A useful distinction is therefore:

local semantic adequacy ≠ cross-domain interpretability

Even with a common EKR foundation, participating domains do not become semantically equivalent. Shared structural, minimal upper-level ontological, identity, and governance commitments can reduce ambiguity and make differences and potential correspondences explicit; remaining differences in domain meaning, granularity, or pragmatic assumptions may still require semantic mediation. This clearer representation can also support compatibility assessment and composable participation.

The EKR capability should therefore create representation readiness for governed cross-boundary use.

Cross-boundary interaction is the stress case emphasised in this article, not the condition that makes a representation an EKR. The concerns that need to be made explicit depend on purpose and scope. A representation maintained within a bounded domain may not need the same cross-domain correspondence mechanisms as one intended to span several domains, while still benefiting from governed identity, authority, lifecycle, and interaction descriptions.

2. One reality, multiple legitimate representations

Enterprise fragmentation is not limited to data silos. It can involve different meanings, structures, identities, levels of granularity, versions, authorities, and operational purposes. A useful way to see the problem is to follow one physical asset through its lifecycle.

Consider a process pump from front-end engineering design (FEED) through detailed design, commissioning, and years of operation. Different parties legitimately describe different aspects of the installation for different purposes:

- FEED and detailed design — process tag, duty, P&ID relationships, equipment specification, and vendor identifiers
- automation and commissioning — controller tags, signals, interlocks, control functions, and possibly an OPC UA information model
- operations and maintenance — runtime history, alarms, an EAM equipment record, component hierarchy, failure modes, and service history

- analytics and AI — time-series context, derived features, model inputs, provenance, and the purpose for which predictions or recommendations are made

These viewpoints are not interchangeable. A process tag may identify a functional position, an automation object may identify a control function, an EAM record may identify a maintainable item, and a vendor serial number identifies a physical unit.

The distinction is not unique to EKR. Process-industry lifecycle standards and handover specifications address closely related distinctions: ISO 15926 covers functions, equipment, installation, operation, maintenance, and replacement across the process-plant lifecycle, while CFIHOS distinguishes tag information describing what equipment does from information about the physical equipment itself. [13] [14] The EKR question is how such identities and relationships remain explicit and governed when several representations and authorities must interact.

The distinction becomes concrete when the pump is replaced. Suppose functional position P-101 is initially occupied by physical pump SN-A and, six years later, by replacement pump SN-B. The process position and much of the automation logic may remain unchanged while the physical asset, vendor record, and maintenance history change. A governed representation must therefore express relationships such as “SN-A occupied P-101 until year 6” and “SN-B occupies P-101 from year 6”, rather than assume that all identifiers are synonyms for one unchanging asset.

Illustrative governed statements for that scope might include:

- P-101 is a functional position in the process design.
- SN-A occupied P-101 until the replacement event; SN-B occupies it afterwards.
- The automation object associated with P-101 describes the control function and is not identical to either physical pump.
- The EAM record is authoritative for the maintained equipment history within its declared scope.
- The representation preserves the provenance, validity interval and source authority of those statements.

In a real implementation, authority would be attached to each statement or representation aspect rather than assumed globally. For example, the authoritative source for the physical replacement event may differ from the authoritative source for the process position or control function.

For a maintenance-handover purpose, an EKR need not copy every source model. It could relate the functional position, automation object, EAM record, and physical equipment identities; preserve which source is authoritative for each aspect; record provenance, version, and validity; and expose only the knowledge needed for the intended interaction. This use of multiple, purpose-specific views aligns with established knowledge-representation principles and architecture viewpoint theory. [2][3]

This leads to the principle of purposeful sufficiency. In this article, a representation is purposefully sufficient when the information required to evaluate or support its declared interaction can be identified and governed without leaving an unresolved representation gap that matters to that purpose. Operationally, this means that the applicable ex-ante readiness questions introduced in Section 8 can be answered positively for the declared use; it does not imply completeness beyond that use. The pump representation required for maintenance handover need not contain everything needed for process simulation, energy optimisation, or control design.

The objective is therefore not to impose one detailed enterprise decomposition. It is to prevent legitimate plurality from turning into uncontrolled ambiguity when interactions cross boundaries.

3. Architecture before technology: separating responsibilities

Modern industrial architectures already provide powerful capabilities through OPC UA, Asset Administration Shell (AAS), Unified Namespace (UNS), APIs, data platforms, ontologies, knowledge graphs, and other mechanisms.

The EKR framing proposes separating architectural responsibilities so that the contribution of each mechanism can be assessed without confusing one responsibility with another.

These are responsibilities, not sequential stages. Governance and security apply across them. Structure means organising data or information according to a schema or form; Represent means making the distinctions, identities, relationships, meaning, authority, and purpose needed for an intended interaction explicit. A system may therefore structure data very well without yet providing the governed representation required for a cross-boundary use.

Used this way, the responsibility lens makes gaps, overlaps and architectural boundaries discussable without prescribing a product stack. Transporting, persisting or exposing information may be necessary, but those functions do not by themselves establish the representation or its governance.

The initial EKR Architectural Guidance workspace has begun to develop this lens into reusable diagnostic questions and patterns rather than prescribe one target architecture.

4. Two ways of constructing governed enterprise representations

The EKR capability must accommodate different authority situations.

When authoritative models or ontologies already exist, the EKR capability may use them as authoritative sources while preserving their scope and authority. Source systems, engineering models, OPC UA information models, enterprise applications, standards-based structures, or externally governed ontologies can remain authoritative in their own domains while selected knowledge is represented for wider use. Depending on purpose and applicable requirements, a local model may itself be an EKR, or it may contribute source-derived knowledge to one.

Where no suitable authoritative representation exists, the EKR capability may maintain operational knowledge directly under EKR authority through authorised knowledge holders, applications or, where governance permits, AI agents operating under explicit provenance, validation, and accountability conditions. Source-derived and EKR-authoritative representations may coexist and may be combined within the same enterprise scope.

There need not be one enterprise-wide “EKR authority”. Authority is scoped to representations and to the statements they contain. Source systems or domain models may remain authoritative for source-derived knowledge. Where knowledge is maintained directly under EKR authority, the accountable owner, steward, or governance role — and the scope within which that authority applies — must be explicit.

This distinction is important because the EKR capability does not require all enterprise knowledge to be copied into, or owned by, one repository. What matters is that the authority applicable to each representation aspect remains explicit. Purpose and scope determine which concerns need to be active in a given representation; crossing boundaries typically brings additional correspondence, provenance, authority, validity, and mediation-readiness requirements into play.

Relevant concerns include:

- authoritative source or responsible knowledge holder
- provenance and derivation
- version and validity
- scope and purpose
- dependencies on other representations
- status of contributions, validation, and publication

Provenance is an established concern in knowledge representation and information governance; W3C PROV-O provides one formal mechanism that can contribute where appropriate. [5]

Lifecycle is equally important. Representations evolve, may be corrected or enriched, may coexist across versions, and may eventually be retired. Operational evidence may inform change, but it does not automatically redefine authority.

5. The shared EKR representation foundation

The Working Definition identifies an explicit shared EKR representation foundation as an architectural requirement. [1] It combines four concerns:

- a shared representation grammar
- minimal upper-level ontological commitments
- identity and relationship principles
- common governance invariants

This is a substantive shared commitment even though EKR rejects the idea of one universal enterprise model. “Minimal” does not mean trivial or easy to agree: it means limiting common commitments to what a declared interoperability scope genuinely requires. The four concerns operate at different levels: grammar provides a formal expression basis; upper-level ontological commitments provide selected cross-cutting semantic anchors; identity and relationship principles govern reference across representations; governance invariants address authority, status, and evolution. The Guidance must determine how little is enough for particular scopes without turning that minimum into a universal domain model.

Shared representation grammar. It provides a common structural and formal basis for expressing representations. It does not impose one domain vocabulary and does not by itself prove semantic equivalence between domains.

Minimal upper-level ontological commitments. They provide a common conceptual frame for a limited set of cross-cutting semantic distinctions needed across representations — for example object, process, role, capability, and selected relationship categories. By stabilising such shared distinctions, they can contribute to representation readiness, including the ability to compare capabilities and assess whether independently developed participants could take part in a common operational composition. The commitment may be thinner than wholesale adoption of any particular top-level ontology: domain ontologies and specialised information models remain free to carry the detailed semantics required by their own communities. Top-level ontology standards provide an established precedent for this role. [6]

No particular representation grammar, upper ontology, or normative implementation profile is selected by this article. Those choices belong to subsequent Guidance development, evidence, and evaluation.

Identity and relationship principles. They make it possible to state what is being represented and how representations refer to or depend on one another.

Common governance invariants. They make authority, contribution, status, and evolution part of the representation discipline rather than external assumptions.

The objective is coherence without semantic uniformity.

This foundation can make heterogeneous representations structurally comparable and create representation readiness for explicit mapping, semantic mediation, and composable participation. It does not itself resolve differences in domain meaning or compose, coordinate, or execute operational interactions. Its role is to support those downstream objectives by making the relevant representations sufficiently explicit, related, and governed.

How much of the foundation should be centralised, distributed, source-resident, or projected through different technologies remains an open implementation question for EKR Architectural Guidance.

6. Building on existing standards and representation approaches

When available, the EKR capability may build on and connect existing representation and integration capabilities.

OPC UA, AAS, UNS architectures, metadata and catalogue platforms, ontologies, knowledge graphs, semantic layers, and data platforms may each contribute to the EKR capability by supplying authoritative models, identifiers, structures, semantics, transport, contextualisation, persistence, governance, or exposure, depending on their actual scope. [7][8] Where those capabilities already satisfy part of the EKR requirement, they should be put to use.

Metadata and context platforms deserve explicit attention because they increasingly extend beyond cataloguing. OpenMetadata, for example, combines technical metadata, lineage, ownership, business glossaries and semantics, classifications, governance, and data contracts, and now presents itself as a context layer for humans and AI agents. [10] For a given scope, a platform of this kind may satisfy substantial EKR responsibilities. The architectural question is which responsibilities it actually fulfils, not whether it belongs to a predefined technology category.

Master Data Management (MDM) is another adjacent discipline. MDM can establish governed identities, ownership, stewardship, and authoritative master or reference records, and may therefore satisfy important EKR responsibilities for a declared scope. [12] EKR does not replace MDM. The additional architectural question is whether the representations needed for the intended interactions — including relevant relationships, provenance, validity, multiple viewpoints, and constraints — are also explicit and governed.

Where those capabilities do not provide the representation required for a declared purpose, EKR-authoritative representations may supply the additional knowledge needed.

The relevant question is therefore:

How can existing and EKR-authoritative representations contribute together to a coherent, governed cross-boundary representation capability?

OMG/EKGF remains an important enterprise-scale reference because it addresses enterprise identity, formal meaning, provenance, governance, contextual plurality, lifecycle, and reuse through an Enterprise Knowledge Graph perspective. [9] Industrial interoperability work should be examined separately. CESMII's current i3X 1.0 specification, for example, defines a common API for contextual manufacturing information across heterogeneous back-end platforms and exposes typed objects, metadata, and relationships, with an associated conformance test suite. [11] Its scope is different from the EKR framing and should be compared responsibility by responsibility rather than treated as equivalent.

Table 1 applies the responsibility lens from Section 3 to a small set of neighbouring approaches. It is illustrative rather than a conformance assessment, maturity ranking, or claim that the same responsibilities are implemented identically in every deployment.

Table 1 — Illustrative application of the EKR responsibility lens to selected neighbouring approaches

Approach	Responsibilities commonly addressed using the Section 3 lens	Scope-dependent questions for EKR assessment
OPC UA	Structure and Represent typed information models, namespaces, and object relationships; Expose operational information. [7]	How are identities related across domains; where do authority, provenance, and lifecycle governance reside; what mappings are needed beyond the adopted information models?
Asset Administration Shell (AAS)	Structure and Represent asset-oriented identifiers, submodels, semantic references and lifecycle information; Persist/Expose digital-twin information. [8]	Which shared commitments span heterogeneous submodels and domains; how are source authority and representation governance handled; where is mediation still required?
Master Data Management (MDM)	Structure/Represent governed master and reference identities; Govern ownership, stewardship, quality, and authoritative records. [12]	Do the governed master records also express the relationships, provenance, validity, viewpoints, and interaction-specific constraints needed beyond master-data scope?
OpenMetadata	Structure/Represent metadata and semantic context; Persist/Expose lineage, ownership, classifications, and contracts; Govern metadata assets. [10]	Which operational identities, capabilities, relationships, and interaction conditions fall inside the governed metadata scope, and which remain elsewhere?
OMG/EKGF	Represent enterprise identity and formal meaning; Persist/Expose knowledge-graph content; Govern provenance, context, lifecycle, and reuse. [9]	How are these responsibilities realised for the industrial scope at hand, and how is fitness for a particular interaction assessed?
CESMII i3X	Transport/Expose contextual manufacturing information through a common API; Structure/Represent typed objects, metadata, and relationships; provide API-level conformance testing. [11]	Which source representations remain authoritative; how are provenance, lifecycle, and cross-domain semantic commitments governed beyond API conformance?

The EKR framing should add architectural discrimination, not merely provide a vocabulary for relabelling existing solutions. For a declared scope, the practical question is whether the architecture can show the required representations, identities and relationships, authority and provenance, lifecycle governance, and fitness for the intended interactions. Where an existing architecture already provides these responsibilities adequately, no additional EKR-specific layer is implied.

Its operational value must therefore be testable. Candidate measures include integration and onboarding effort; effort to establish and maintain semantic correspondences; change-impact effort; time to assess compatibility; reuse of governed representations; and ambiguity or rework attributable to unclear identity, authority, provenance, or purpose. Implementation cases should determine which measures are meaningful and whether the EKR framing produces a material improvement for the scope being examined.

For the maintenance-handover example, a testable hypothesis is that making the relevant identities, relationships, authority, provenance, and validity explicit before handover reduces unresolved identity/authority ambiguities and later reconciliation effort relative to the project's existing handover baseline. If those outcomes do not improve, the EKR treatment has not demonstrated operational value for that scope.

Conversely, if the same outcomes can be achieved with no useful architectural distinction, no relevant reduction in ambiguity or effort, and no improvement in the intended downstream interactions, that is evidence that additional EKR structures may not add value in that scope.

7. Governing representation assets

Once representations become enterprise assets used across boundaries, governance applies to more than data flows or storage locations.

A practical difficulty is that representation decisions and interpretability costs often fall on different actors at different times. On a large automation programme, identities and model structures may be fixed during FEED and detailed design by project organisations, suppliers, and integrators, while the cost of ambiguity may emerge years later for the operator at handover, integration, or change. For important scopes, representation requirements may therefore need to appear as project deliverables and acceptance criteria, supported by reusable standards or profiles where appropriate and then sustained under the operator's own governance. The appropriate enforcement point — contract, standard, operator governance, or a combination — is context-dependent.

This also creates an economic question. The party that pays to make identity, relationships, provenance, and authority explicit may not be the party that later benefits from lower integration or ambiguity costs. Project suppliers may bear part of the representation cost while operators, maintainers, integrators, or future application teams receive much of the downstream benefit. That asymmetry is one reason representation requirements may need to be expressed in project deliverables and acceptance criteria rather than left to voluntary late-stage cleanup. The hypothesis to test is that this additional effort reduces later onboarding, change-impact, and ambiguity costs enough to justify the governance burden; EKR should not assume that benefit in advance.

Relevant concerns include authority, provenance, versioning, validity, contribution and publication rights, access control, classification, compartmentalisation, auditability, and lifecycle management.

Security and sovereignty also apply at representation level. A representation can reveal organisational structures, equipment capabilities, dependencies, operating constraints, semantic correspondences, or interaction conditions even when no live process value is exposed.

Relevant architectural questions therefore include:

- who can access which representation
- who can contribute, validate, publish, or retire it
- what the representation reveals and under which permitted uses
- how access and classification are separated across organisational or security domains
- whether provenance, versions, and changes remain auditable
- how the representation can be exported, migrated, or reconstructed

- how source authority remains visible after knowledge is exposed elsewhere

These are representation-governance concerns that the EKR capability must address, but they do not yet constitute a complete governance operating model. Future guidance must clarify which controls belong inside the representation capability and which are supplied by existing identity, security, policy, audit, or organisational-governance mechanisms.

8. From governed representation to interoperability and composability readiness

A governed representation can make operational participants explicit in terms of capabilities, requirements, inputs and outputs, constraints, dependencies, applicable contexts, authorities, and interaction conditions.

Descriptions of this kind can support discovery, onboarding, compatibility assessment, substitution analysis, and preparation for composable participation.

Composable participation means that a participant — an asset, application, service, organisation, or agent — is represented well enough for another participant or downstream mechanism to discover it, understand the capability it offers, assess relevant constraints and compatibility, and determine whether it can take part in an intended interaction. This can matter within a domain as well as across domains.

Operational composability is the broader objective of being able to combine, substitute, or reconfigure such participants as operational needs change — for example when integrating a new system, replacing an equipment capability, migrating a function between implementations, onboarding a partner, or adapting an operational process. It does not imply that those changes are automatic: the EKR contribution is to make the relevant capabilities, constraints, and relationships explicit enough for downstream mechanisms and governance to assess and carry them out.

The controlled EKR boundary is:

Operational composability is an EKR-enabled objective supported by governed representations and interaction/capability descriptions; composition, coordination and execution remain downstream capabilities rather than intrinsic EKR functions.

The EKR capability can expose governed representations and interaction descriptions containing the elements needed to assess possible interactions. Semantic mediation may establish correspondences where domain meanings differ. Other downstream mechanisms may then perform compatibility decisions, authorisation, negotiation, coordination, or execution.

The EKR capability therefore creates representation readiness: it makes the required knowledge explicit, identifiable, traceable, governed, and sufficiently structured for downstream capabilities to operate on it. Readiness should be assessed against a declared downstream use before that use is attempted, then tested by the downstream outcome.

Candidate ex-ante readiness questions include:

- Are the participants, objects, and relevant relationships unambiguously identifiable for the declared use?
- Are authority, provenance, version, and validity explicit where they affect the interaction?
- Are the capabilities, constraints, or contextual conditions needed for the interaction represented?
- Are known semantic or granularity differences exposed so that mediation can address them explicitly?
- Can the representation be evaluated and evolved without losing the source authority on which it depends?

For the pump handover, these checks ask whether the functional position and physical equipment are distinguished; whether current and historical occupancy relationships are explicit; whether source

authority, provenance, version, and validity are known; and whether the relationships needed by maintenance are represented. A downstream mediation or onboarding mechanism can then test whether those preconditions were sufficient. If a failed interaction can be traced to a missing or ambiguous representation element, that is evidence that the representation was not ready for that purpose; failures caused by downstream mechanism defects or other causes must be distinguished.

9. Open architectural guidance is now underway

The Working Definition and this first article provide the conceptual starting point for EKR Architectural Guidance. The public repository is the open development workspace for evidence, implementation experience, challenges, counterexamples, and Draft Guidance.

The current V0.1 backlog asks, among other things, what minimum representation information is required, what belongs in the shared foundation, how source authority and provenance should be preserved, how conformance should be distinguished from fitness for purpose, where EKR responsibilities stop, and what implementation cases would constitute useful evidence.

The complete ten-question backlog is reproduced in Appendix A as a publication-time snapshot; the repository remains the living source as the Guidance evolves.

Practitioners, standards communities, researchers, technology providers, and industrial users are invited to challenge the framing, identify existing approaches that already satisfy these responsibilities, and contribute cases in which the EKR lens does — or does not — provide useful architectural discrimination.

10. Conclusion

The value 85 is easy to transport. The knowledge required to interpret and use it reliably is not carried by the value itself.

That knowledge must remain identifiable, governed, and traceable, and must be prepared for the semantic mediation or compatibility assessment that may be required when different domain representations meet. The pump example shows the same issue over a lifecycle: the cost of a missing identity, relationship, or authority statement may surface long after the data and models were first created.

Within a domain, modelling, information models, and ontologies — when available — can establish the shared concepts and conventions needed by their participants. The EKR capability can use those models as authoritative sources where appropriate, while also allowing operational knowledge to be represented directly under EKR authority where no suitable authoritative model exists.

A shared EKR foundation provides common structural, selected upper-level ontological, identity, and governance commitments that make heterogeneous representations more explicit and relatable. Where domain meanings still differ, semantic mediation may remain necessary.

The objective is a coherent capability for managing multiple purposeful representations across systems, domains, organisations, and time while preserving legitimate local authority.

The next stage is therefore both architectural and empirical. EKR Architectural Guidance is being developed through decision questions, patterns, governance principles, evaluation criteria, and implementation experience, while worked cases must test whether representation readiness can be demonstrated and whether the resulting benefits justify the effort required.

The present article draws mainly on industrial examples; future Guidance and implementation cases should test whether the same representation responsibilities remain useful in non-industrial enterprise settings.

That development is open to challenge through the public EKR repository. EKR should mature not by making every existing architecture fit the framing, but by confronting cases in which its distinctions help, cases in which they need refinement, and cases in which they add no useful value.

Public development resources

- Public EKR repository: <https://github.com/jvieille/ekr>
- EKR Architectural Guidance workspace: <https://github.com/jvieille/ekr/tree/main/guidance>
- Contribution guide and repository governance:
<https://github.com/jvieille/ekr/blob/main/CONTRIBUTING.md>

Acknowledgements

The author gratefully acknowledges William H. Bosler (Independent Consultant), François Jantet (Orano), Bindia Ouali (Safran), Hervé Panetto (University of Lorraine, TELECOM Nancy, CNRS), Stéphane Puydarrieux (Orano), and Swarandeeep Singh (ABB), for their positive feedback, thoughtful review, and critical comments on successive drafts of this article. Their questions and challenges helped clarify the practical scope of EKR, its relationship to existing approaches, the need for testable operational evidence, and the governance implications of representation across the asset lifecycle. The final arguments and any remaining errors are the author's responsibility.

AI-assisted drafting and review disclosure

Generative AI tools were used during the development of this article to support editorial drafting, structured critique, consistency checking, and review of alternative formulations. OpenAI ChatGPT was used throughout the drafting and revision process; Anthropic Claude and DeepSeek were also used for independent critical review of later drafts. All suggestions were evaluated by the author, who selected the sources, determined the architectural positions, and remains solely responsible for the arguments and final text.

References

- [1] Vieille, Jean. Enterprise Knowledge Representation — Working Definition V0.1. 14 August 2026. DOI: 10.5281/zenodo.21914714.
- [2] Davis, Randall, Howard Shrobe, and Peter Szolovits. "What Is a Knowledge Representation?" *AI Magazine* 14, no. 1 (1993): 17–33. DOI: 10.1609/aimag.v14i1.1029.
- [3] ISO/IEC/IEEE 42010:2022. Software, systems and enterprise — Architecture description. 2nd ed., November 2022.
- [4] W3C OWL Working Group. OWL 2 Web Ontology Language: Document Overview (Second Edition). W3C Recommendation, 11 December 2012.
- [5] Lebo, Timothy, Satya Sahoo, and Deborah McGuinness, eds. PROV-O: The PROV Ontology. W3C Recommendation, 30 April 2013.
- [6] ISO/IEC 21838-1:2021. Information technology — Top-level ontologies (TLO) — Part 1: Requirements. 1st ed., August 2021.
- [7] OPC Foundation. OPC 10000-1, OPC Unified Architecture — Part 1: Overview and Concepts; and OPC 10000-5, OPC Unified Architecture — Part 5: Information Model. Release 1.05.06, 22 October 2025.
- [8] Industrial Digital Twin Association. Specification of the Asset Administration Shell Part 1: Metamodel, IDTA-01001, Version 3.2, July 2026. DOI: 10.62628/IDTA.01001-3-2.
- [9] Object Management Group. Enterprise Knowledge Graph Platform Task Force; Enterprise Knowledge Graph Forum, EKG Principles. <https://www.omg.org/ekg/> ; <https://principles.ekgf.org/principle/> . Accessed 21 September 2026.
- [10] OpenMetadata. OpenMetadata — Open Context Layer; Data Contracts documentation. <https://open-metadata.org/> ; <https://docs.open-metadata.org/v1.11.x/how-to-guides/data-contracts> . Accessed 21 September 2026.

[11] CESMII. i3X — Industrial Information Interoperability eXchange, Version 1.0 Specification and Conformance Test Suite. <https://github.com/cesmii/i3X> ; <https://www.i3x.dev/> . Accessed 21 September 2026.

[12] DAMA International. DAMA-DMBOK2 Revised Edition. 2024. <https://dama.org/dama-dmbok2-revised-edition-faqs/> . Accessed 21 September 2026.

[13] ISO 15926-2:2003. Industrial automation systems and integration — Integration of life-cycle data for process plants including oil and gas production facilities — Part 2: Data model. Current edition reviewed and confirmed in 2023. <https://www.iso.org/standard/29557.html> .

[14] CFIHOS. CFIHOS Standards V2.0, including Scope and Procedures, Reference Data Library and Data Model. <https://www.jip36-cfihos.org/cfihos-standards/> . Accessed 21 September 2026.

Appendix A — Open-question backlog snapshot

The following questions reproduce the EKR Architectural Guidance V0.1 backlog as it stands at the time of this article revision. The repository remains the living source.

- What is the minimum representation information required before a cross-boundary interaction can be considered adequately governed?
- Which commitments must belong to the shared EKR representation foundation, and which should remain domain-specific?
- When should knowledge remain source-resident and be projected, versus be contributed directly under EKR authority?
- How should authority, provenance, version and validity be represented when multiple sources describe different aspects of the same operational reality?
- How should conformance to an EKR foundation be distinguished from fitness of a representation for a particular operational purpose?
- How should multiple viewpoints and granularities coexist without collapsing into one enterprise decomposition?
- Where should EKR responsibilities stop and semantic mediation, compatibility assessment, coordination or execution begin?
- What representation-level security, classification and sovereignty controls are required beyond data access control?
- How can existing OPC UA, Asset Administration Shell (AAS), Unified Namespace (UNS)/DataOps, ontology, knowledge-graph and data-platform capabilities be assessed by architectural responsibility rather than by product category?
- What should worked examples or implementation cases demonstrate to provide useful evidence for the first Guidance release, and what would make such examples or cases misleading?