A Global Cooperation Model for Sovereign Industrial AI

European dataspaces and open standards for verifiable evidence across platforms

Version 1.0

Dr. Carsten Stöcker

Spherity GmbH

Prepared: 1 November 2026



Except where otherwise noted, this work is licensed under the Creative Commons Attribution 4.0 International License (CC BY 4.0). © 2026 Dr. Carsten Stöcker, Spherity GmbH.

Abstract

Industrial AI needs infrastructure that connects evidence across companies and jurisdictions at operational speed. European dataspaces are a first step in participant-controlled cooperation, not a demonstrated infrastructure foundation for IPCEI-AI. Dataspace concepts and policy precede today's LLM, frontier and agentic AI wave; existing governance and exchange mechanisms do not establish fitness for high throughput, low latency and large data volumes. We argue for **fast-track modernisation** that meets workload-specific performance requirements without relaxing access control, verifiability or participant authority. Substantial engineering and validation remain.

Palantir's operational ontology provides a demanding architectural reference for connecting data, meaning, logic and governed actions. Our proposed counter-model extends that operational ambition across independently governed implementations, with credible provider choice. Its foundation is **access-controlled, verifiable evidence graphs built on open international standards and industry-agreed semantics**. These graphs connect decision-relevant claims and provenance while bulk data can remain in specialist systems.

We combine primary-source review, platform economics, comparable architecture patterns and a proposed industrial experiment. The resulting agenda joins efficient data movement, task-specific graph views, scoped authorization and maintained evidence status. Europe can lead this modernisation through global coordination with US innovators, partners across Asia-Pacific including China, and other participant-controlled ecosystems. Joint vocabulary governance, implementation profiles and reciprocal testing must turn shared principles into working interoperability. This is a collaborative blueprint for global digital sovereignty and a competitive multipolar economy, not a claim of production readiness. Performance, assurance, resilience and economic gains remain outcomes to demonstrate together. Accelerated investment should fund redesign and testing where needed, rather than assume that existing dataspace implementations already satisfy AI-first requirements.

Keywords. industrial AI; digital sovereignty; digital diplomacy; Palantir; dataspaces; platform economics; semantic interoperability; verifiable evidence; international cooperation

Contents

1 Introduction **4**

2 Research design and scope **5**

3 Platform economics and the industrial meta platform **6**

3.1 Why integrated platforms create value 6

3.2 Cooperation beyond the platform boundary 7

3.3 Avoiding a new intermediary monopoly 7

4 Meaning control and the limits of openness **8**

4.1 Separate the dimensions of centralization 8

4.2 Vocabulary agreement is different from a reasoning assumption 8

4.3 The ontology engine as an ecosystem capability 9

5 Pattern P1 The operational ontology platform **11**

5.1 Common functions and the ontology engine 11

5.2 Ingestion is followed by modeling and graph construction 11

5.3 AI uses governed tools and actions 13

5.4 Multi company operation and the portability question 13

6 Pattern P2 Participant controlled industrial cooperation **13**

6.1 Preserve local capabilities and change the cooperation boundary 13

6.2 Shared meaning discovery and trust 15

6.3 Admit evidence before using it for assurance 16

6.4 Reason over the authorized graph and act under local authority 16

6.5 Policy interoperability and the secure semantic supply chain 17

7 Comparative analysis of control and dependence **17**

7.1 A counter model with the same operational obligations 17

7.2 Documented risks and changing provider behavior 19

7.3 Geopolitical exposure and system level dependency 19

7.4 Evidence and assurance as portable industrial capabilities 20

8 A humanoid robot example linking evidence to scoped decisions **22**

9 An industrial example and a falsifiable comparison **26**

10 Existing foundations and unresolved implementation work **27**

10.1 European coordination with international participation 27
10.2 Fast-track modernisation for AI-first operations 27
10.3 Standards do not compose themselves 28

11 Geopolitical context and digital diplomacy 28
11.1 A global counter model based on reciprocal agency 28
11.2 European leadership through shared implementation 29
11.3 UNECE traceability and transparent standards governance 29
11.4 Asia Pacific partnership and regional choice 30
11.5 US innovation and global web standards 31
11.6 Digital diplomacy through bounded cooperation agreements 32

12 An international industrial and procurement strategy 34
12.1 Fund cooperation as an operational capability 34
12.2 Require an exit that preserves useful behavior 34
12.3 Govern the commons without centralizing the market 35

13 Scenarios limitations and the research agenda 35

14 Conclusion 36

15 Data availability and disclosures 37

A Common blocks and detailed local flows 37

B Federation exchange and dependency contracts 38

C Candidate components and planned experiment 39

References 42

1 Introduction

A manufacturer asks an AI assistant whether a replacement valve can enter a demonstration water loop. A supplier identifies the component, an independent laboratory reports a test result, and the manufacturer holds its acceptance criteria and inventory state. The assistant must establish that the records describe the same valve, interpret the test conditions, check the laboratory's accepted authority and determine whether the requester may initiate the release workflow. A fluent answer is insufficient. The manufacturer needs an accountable connection between evidence and action. In a global supply chain, the organizations may use different platforms and operate under different laws.

Palantir's operational ontology is a powerful response to the integration problem. Its documented architecture joins data, logic, actions and security through a Language, Engine and Toolchain. This combination makes the ontology more than a dictionary or a graph visualization: it becomes part of how an organization operates. That architectural contribution deserves explicit recognition. The proposition that it alone explains Palantir's commercial success would require separate business evidence and is not advanced here. [1]

The next architectural opportunity is to make comparable cooperation available across independently governed environments. A Japanese supplier should be able to serve customers using different systems. A US software provider should be able to implement an agreed interface and compete for their business. A Chinese laboratory's evidence should be assessable against a declared scope and applicable rules without assuming either universal acceptance or categorical exclusion by nationality. These are proposed capabilities, not claims about an existing global recognition arrangement. The shared objective is useful cooperation with accountable local decisions.

Europe has substantial experience in organizing this effort. Its Common European Data Spaces programme combines sector participation, infrastructure and governance to support controlled data sharing. The underlying dataspace research dates to 2005, while the European data strategy set out Common European Data Spaces in February 2020. These developments precede today's LLM, frontier and agentic AI wave, although they do not predate AI itself. This history explains why new workload requirements need explicit assessment; it does not establish that every existing implementation is obsolete or slow. We describe Europe's contribution as pioneering industrial coordination rather than exclusive invention of dataspaces, knowledge graphs or open standards. [2] [3]

Dataspaces are therefore a first step, not evidence of readiness for AI-first operations or IPCEI-AI. We argue for **fast-track modernisation**: assess, redesign and where necessary replace components against declared throughput, latency, data-volume and trust requirements. The prior evidence-graphs paper motivates separating high-frequency data processing from the semantic context and evidence needed to assess its outputs; it presents a design proposal, not a comparative performance benchmark. The next phase must combine this engineering work with international coordination among ecosystems pursuing participant-controlled infrastructure. A model confined to Europe would not meet the needs of global production and AI value chains. [4]

We therefore position **participant-controlled industrial cooperation as a global counter-model to platform dependence, with Europe as a leading convener and implementation partner**. It is Europe's constructive strategic answer to Palantir-centred integration and closed Big Tech ecosystems: preserve operational depth while making cooperation possible across platforms. Palantir and other capable providers can participate in such an arrangement. The distinction concerns dependence and authority, rather than an opposition between European and American innovation.

The EU's 2025 International Digital Strategy already joins competitiveness with international partnerships and proposes technical exchanges and interoperability pilots. The Commission's cloud-sovereignty framework separately considers legal, operational, technological and supply-chain dependencies. These documents provide policy context for a cooperation model in which sovereignty means the effective capacity to choose partners, govern permitted uses and preserve essential operations. We propose making that capacity an industrial objective of digital diplomacy; we do not claim that P2 is an adopted EU architecture. [5] [6]

Shea, Smith and Stöcker's *Decentralized Identity as a Meta-platform* offers a conceptual precedent for participant-controlled cooperation across platforms. We extend its question from identity to industrial semantics, evidence and action. This is an architectural and economic extension of prior co-authored work, not independent empirical validation of its network-value claims. [7]

Our research question is: **Under what architectural and governance conditions can industrial AI preserve useful meaning, verifiable evidence and authorized behavior across independent platforms and jurisdictions, supporting cooperation without avoidable dependence on a single commercial or technical intermediary?** The anchor is **access-controlled, verifiable evidence graphs built on open international standards and industry-agreed semantics**. We examine how this infrastructure could strengthen international trust and a competitive multipolar digital economy, while keeping documentary findings, design proposals and expected benefits distinct.

2 Research design and scope

This article combines targeted documentary review, comparative architectural synthesis and a proposed matched-task experiment. The evidence cutoff is 1 November 2026. The source corpus includes Palantir documentation, normative specifications, implementation releases, competition-authority findings, European policy and legal instruments, first-party ecosystem reports and selected research literature. Sources were selected for a specific capability, dependency mechanism, conceptual distinction or maturity claim. This is not a systematic review of all enterprise AI systems.

Four kinds of statement remain distinct. A documentary finding reports what an identified source establishes within its scope. An architectural proposal specifies desired behavior. A strategic judgment recommends a course of action from stated premises. An analytical scenario describes a possible development without assigning an empirically estimated probability. No authenticated Palantir deployment, composed federation, security exercise or performance benchmark was executed. The article does not rank products by unmeasured security, cost or intelligence.

The frozen architecture baselines are P1 version 0.5 and P2 version 0.4. They share twelve functional blocks, twenty-five corresponding local flow groups and eleven assessment fields per block. P2 adds twelve federation exchange groups. These are properties of the analytical representation, not counts of deployed services. The present manuscript revision integrates the patterns in the main text and moves detailed flow, exchange and test specifications to appendices and versioned supplements.

Palantir is the documented example of an integrated operational-ontology platform. It is not a proxy for every US technology company, and this article does not establish that Palantir is a monopoly. Hyperscaler cloud markets supply separate evidence about infrastructure concentration and dependence. The article connects these layers through a dependency analysis

rather than assuming they share one architecture or market position. Similarly, P2 is an author-proposed global cooperation model informed by European dataspaces and international standards, not a claim that every data space implements it. References to multipolarity describe multiple effective centres of participation and innovation; they do not imply a uniform political alliance or governance regime.

The contributions are a theory-informed account of industrial platform dependence, comparable architectural patterns, explicit portability requirements, a digital-diplomacy agenda and a protocol that can falsify claims of preserved behavior. Platform-owner and programme-owner reports are attributed rather than treated as independent audits. Publication and report dates are distinguished from access dates; rolling documentation is not silently equated with a pinned software release.

The source register records retrieval limits. The meta-platform analysis used its complete Markdown version; equivalence to the linked final PDF was not established. Van Hove's methodological article was assessed through its institutional abstract and bibliographic record. The Data Act and Data Union discussion uses the Commission's official explanations where linked EUR-Lex texts were inaccessible. The Recommendation 46 and 49 discussion relies on inspected official indexed descriptions and decision text because direct UNECE documents returned access errors; a full-document read is not claimed. Chinese official texts are paraphrased in author translation. These sources support the bounded claims made here, not unseen detailed analyses or instrument-level legal advice.

3 Platform economics and the industrial meta platform

3.1 Why integrated platforms create value

An integrated platform can reduce the number of interfaces an organization must reconcile. Its developer tools, identity controls, operational model, applications and deployment procedures can be designed together. Engineers can reuse objects and actions across use cases instead of rebuilding an integration for every application. Customers may rationally prefer this coordination, even when alternatives exist. The architectural counter-model has to account for these benefits rather than treating integration itself as a defect.

Established systems-economics research helps explain the trade-off. Katz and Shapiro examine compatibility, network effects and competition among systems. Compatibility can broaden component choice, while coordinated systems and network sponsorship can support investment and adoption. Achieving compatibility also entails costs and can constrain design choices. Their analysis provides a basis for asking which interfaces should be standardized and which should remain open to differentiated implementation. It does not prove that a specific contemporary federation will outperform an integrated platform. [8]

For industrial AI, network value can arise from reusable domain models, available integrations, trusted counterparties and a growing set of operational applications. The same accumulated investments can make migration difficult. A customer may be able to export records while losing application behavior, policy mappings or a usable explanation of how earlier decisions were made. This article calls that risk **operational dependence**. It concerns the cost and feasibility of preserving required behavior, not merely possession of a data copy.

3.2 Cooperation beyond the platform boundary

The meta-platform paper argues that identity infrastructure can support participant-controlled value transfer across other platforms. Its distinction between aggregation and cooperation concerns who can carry relationships and capabilities between networks. It recognizes costs of portability, gateway intermediaries and usability, and acknowledges circumstances in which integration or acquisition remains attractive. Its illustrative network-value expressions should be read as conditional models; a quadratic expression is not exponential growth. The paper does not establish a measured law of industrial-AI value. [7]

Smith's originating network-scaling essay supplies part of that conceptual argument, while Van Hove's methodological discussion cautions that tests of network-value formulas depend on the variables and assumptions selected. We therefore do not forecast economic returns from participant counts or apply a network-growth formula to European GDP. The relevant question is whether useful cross-platform cooperation adds more value than the semantic, verification, coordination and governance work needed to sustain it. [9] [10]

Our extension is to treat identity as necessary but insufficient. An identified organization still needs to understand a claim, determine the issuer's authority, respect its conditions of use and decide what operation may follow. An industrial meta-platform would make these relationships portable through maintained exchange contracts. It would allow competing platforms to supply sophisticated local capabilities while reducing the need to integrate every participant into the same operational environment.

This is a proposition about the organization of markets as well as software. A supplier that can reuse an accepted test assertion with several customers may avoid repeated onboarding. A software provider that can implement an agreed evidence interface may compete for an existing customer without reconstructing every supplier relationship. A regulator may inspect an attributable decision package without requiring all regulated firms to use one vendor. These are anticipated mechanisms, not measured savings from the present study.

3.3 Avoiding a new intermediary monopoly

Cooperation can reproduce the dependencies it is intended to reduce. A mandatory gateway operator, proprietary vocabulary service or indispensable trust registry can become a new concentration point. Open-source availability does not eliminate the cost of operating an alternative. A participant may nominally own its data while depending on a service provider for keys, status information, mappings or audit history. The proposed counter-model therefore requires practical replacement and continuity arrangements at those interfaces.

Decentralization should be applied selectively. Some coordination is desirable: a shared identifier scheme, a sector profile or an accepted accreditation authority can make exchange useful. Requiring a participant to replace a legally mandated authority would be incoherent. The design question is which authority is justified by the task and which commercial or technical dependency is incidental. Shared governance must make that distinction explicit, publish its decision procedures and provide appropriate challenge and change processes.

The economic evaluation must compare total operating arrangements. Federation adds profile maintenance, dispute resolution, evidence checking, distributed recovery and conformance work. Integrated platforms also incur integration, customization, governance and exit costs. A fair comparison measures these costs against comparable useful outcomes over a declared period. Provider substitution is a valuable test because it exposes dependencies that remain invisible while all parties continue using the original configuration.

Aggregate gains do not guarantee that every participant benefits. A manufacturer may save integration effort while a small supplier bears new evidence-production costs, a laboratory assumes additional liability, or a gateway operator lacks sustainable revenue. Evaluation should therefore report costs, benefits and bargaining conditions separately for these roles. Participation, payment and governance arrangements must make cooperation worthwhile for the parties whose continuing contribution it requires.

4 Meaning control and the limits of openness

4.1 Separate the dimensions of centralization

The term centralized is ambiguous. Data can remain in several locations while one platform governs the usable model and application interfaces. Conversely, several firms can use a managed shared service under governance that preserves independent authority and credible exit. This article distinguishes physical data placement, runtime coordination, semantic governance, identity and policy control, commercial dependence, and legal authority. A geography label determines none of these dimensions on its own.

The P1/P2 comparison concerns the governance of operational cooperation. P1 integrates capabilities around a configured platform environment. P2 makes selected boundary contracts usable by independently governed implementations. An integrated platform can be a P2 participant. A federation can still depend heavily on US cloud services or model providers. This distinction prevents a misleading promise that an interoperable knowledge graph by itself resolves semiconductor supply, cloud concentration, model access or jurisdictional exposure.

Strategic autonomy is used here in an operational sense: the ability to continue essential functions, exercise lawful authority, scrutinize dependencies and change a replaceable provider at an acceptable cost. It is compatible with international trade and technical interdependence. The relevant procurement question is whether an alternative can actually preserve the required function, not whether the system is described as sovereign or decentralized.

4.2 Vocabulary agreement is different from a reasoning assumption

Knowledge-graph research distinguishes graph structures, schema, identity, context and deductive or inductive techniques. Research on language models and graphs likewise separates graph-supported models, model-assisted graph construction and reciprocal combinations. These distinctions require separate architectural responsibilities for preparation, modeling, graph building, reasoning and action. A graph representation alone supplies no universal AI reasoning regime. [11] [12]

Under the open-world assumption explained by OWL, an absent statement is not automatically false. That is useful when knowledge is distributed and incomplete. It does not imply public data, open-source implementation or international consensus. An operational gate can still require a defined set of acceptable premises in an authorized snapshot. SHACL can validate such a supplied graph against selected constraints, including closed shapes, without establishing the truth of its measurements. [13] [14]

Each implementation profile must declare the reasoning regime it actually uses. Conflicting claims remain separately attributed; combining inconsistent axioms indiscriminately into a single classical theory is not the proposed design. Validation, inference, statistical prediction and an LLM's explanation have different acceptance conditions and should remain distinguishable in the decision record.

Semantic interoperability requires representation machinery, domain agreement and a precise implementation profile. RDF and JSON-LD provide exchangeable graph structures and identifiers. A vocabulary supplies definitions; a profile selects terms, units, identifiers, constraints and extension rules. GS1's Core Business Vocabulary illustrates agreement within a defined event domain. Its existence does not supply every industrial concept or remove the need for governed mappings. [15] [16] [17]

For example, a supplier's released status may mean that internal inspection is complete, while a manufacturer's released status authorizes use in a particular installation. Matching the strings would create a false equivalence. A mapping needs an owner, source and target versions, scope, transformation and known information loss. AI can propose mappings, but their promotion into operational use needs validation and accountable approval.

Franklin, Halevy and Maier's earlier dataspace research distinguishes administrative proximity from semantic integration. This helps explain why cross-company access and shared meaning are related but separate problems. Contemporary data-space programmes add their own governance and exchange protocols; they should not be treated as identical to that earlier research abstraction. [3]

4.3 The ontology engine as an ecosystem capability

Figure 1 makes the architectural progression explicit. P1 integrates an ontology model, graph construction and operational services within a configured platform governance scope. A company or agency is a useful focal example, while Palantir also documents multi-organization ecosystems and industry-standard or participant-defined common ontologies. Enterprise focus is therefore a presentation of the governance pattern, not a single-company product limit. Nor is a configured model evidence that every Palantir workload applies the formal closed-world assumption. [1] [18]

P2 extends the ontology capability across independently governed participants. Each retains its own model, graph builder and runtime; agreed semantics, verifiable assertions and policy-bound exchange connect their work. Shared Meaning, Discovery and Trust support every participant. The next step is an interoperable ecosystem of ontology engines and authorized AI tools, with accountable local decisions. A federated query or disclosed assertion can extend the evidence available to a task without moving the entire graph into one platform.

The open-world focus matters when evidence is distributed: a laboratory result that is unavailable, undisclosed or not yet linked must not be treated as proof that no result exists. A selected reasoning profile keeps that uncertainty visible. A consequential operation can still require a complete set of accepted premises and return HOLD when they are missing. The intended advance combines cross-company meaning, evidence and authority; open-world logic alone does not supply interoperability or permission. [13] [14]

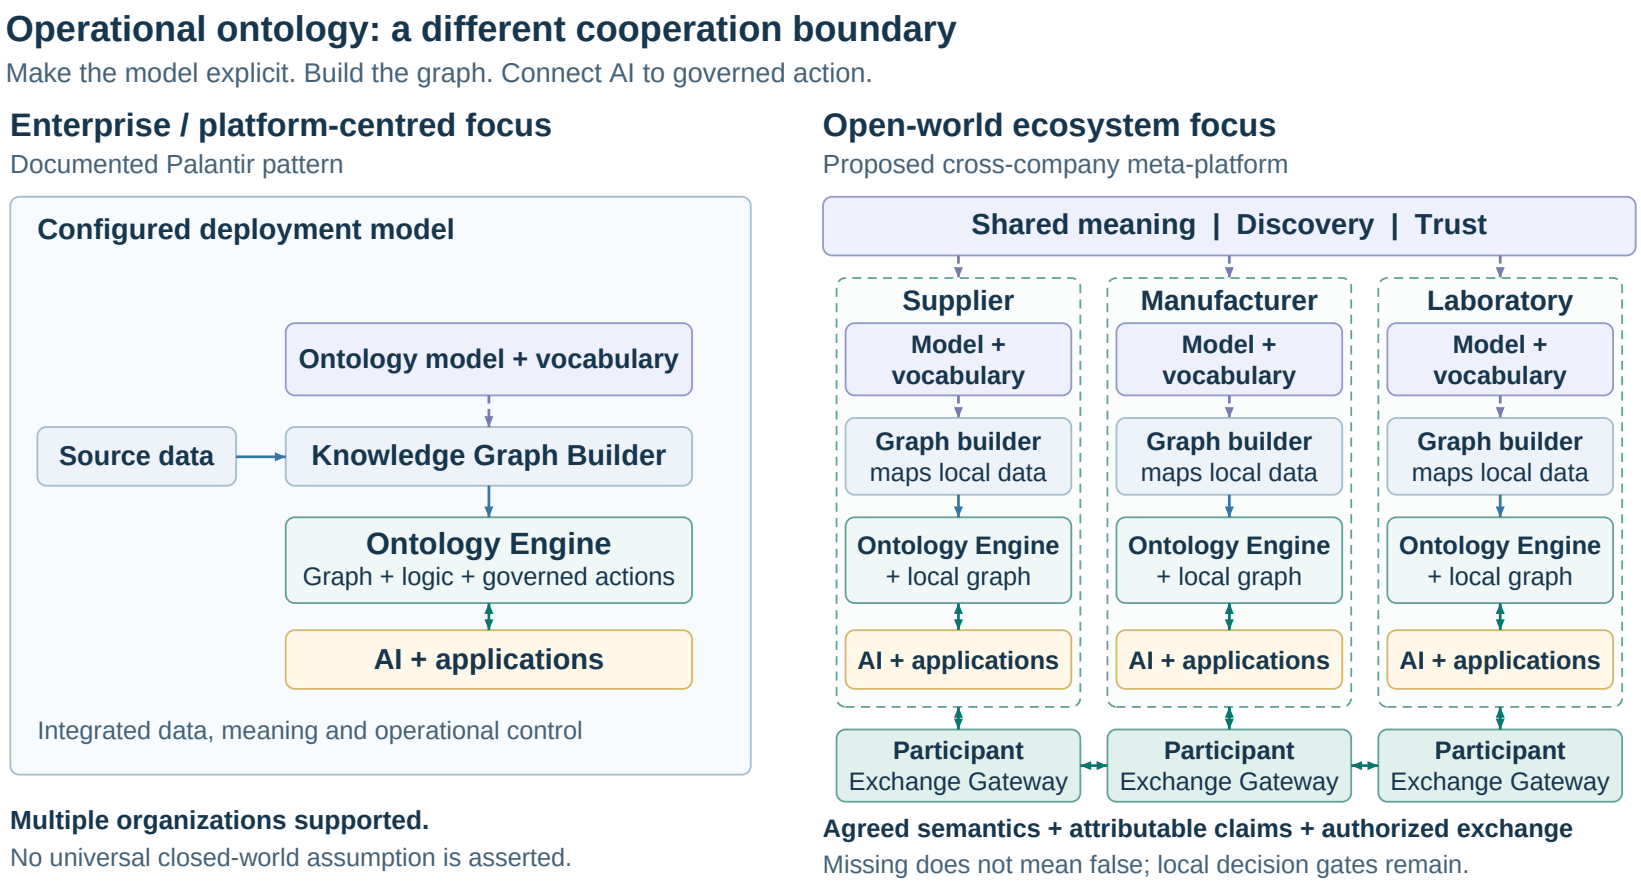


Figure 1: Conceptual comparison of ontology capability within a platform governance scope and across independently governed participants. Both require an ontology model, graph construction and an operational runtime; local ingestion and control loops are abbreviated. The proposed ecosystem connects independently governed participants through shared meaning, discovery and trust, scoped exchange and cross-company AI. Palantir supports multiple organizations; the left-hand focus does not assert a single-company limit or universal closed-world reasoning. On the right, open-world treatment of incomplete knowledge coexists with local authorization and bounded decision gates. This is a proposed extension, not a measured superiority result.

5 Pattern P1 The operational ontology platform

5.1 Common functions and the ontology engine

The twelve common blocks describe responsibilities rather than twelve separate products. B01 accesses sources, B02 prepares them and B03 supplies the data substrate. B04 defines the ontology model and vocabulary. B10, the Knowledge Graph Builder, maps and maintains connected instances. B05 serves the operational graph; B06 executes logic and models; B07 orchestrates AI; B08 executes authorized actions; B09 supplies applications. G01 governance and G02 lifecycle apply throughout. Appendix A gives the complete block and flow mapping.

Palantir uses the term Ontology Engine within its own architecture. For comparison, the neutral ontology runtime groups graph reads, executable logic and actions under governance and lifecycle control. The Knowledge Graph Builder is our explicit functional abstraction. It overlaps the engine's write and indexing path; it is not asserted to be a separately named Palantir product. Figure 2 makes the model, builder, runtime and their principal flows visible. [1] [19]

5.2 Ingestion is followed by modeling and graph construction

Structured records and files enter through configured connectors or pipeline inputs. Palantir documents Excel access and document processing, including extraction from PDFs and images. Format support must be assessed against the particular route and release; the inspected Excel connector is beta. A generic list of office file types does not establish that every Word or PowerPoint file receives reliable semantic interpretation automatically. Reading bytes, extracting a value and establishing business meaning are separate tasks. [20] [21] [22] [23]

The principal path is B01 → B02 → B03 → B10 → B05. Preparation extracts, normalizes and reconciles source material. The substrate retains suitable data, media or remote references. The ontology model provides definitions to the builder and runtime. Model-aware mappings create and maintain object instances, properties and relationships. Palantir's multimodal architecture also includes virtual access and external compute, so the logical pattern does not require copying every original asset into a single store. [24] [25]

In the valve case, a database identifier, spreadsheet value and laboratory report must first be associated with the correct component and test attempt. Object and link configuration make those associations operational. The backend documentation describes the Object Data Funnel indexing source changes and action edits into object databases, while object services support reads over the populated representation. These mechanisms explain graph construction without assuming an RDF triple-store implementation or native OWL entailment. [19] [26] [27] [28]

The evidence does not establish one universal formal closed-world inference regime for Palantir. Proprietary implementation and enterprise governance are insufficient grounds for that classification. The contrast with P2 concerns independently governed cooperation and the required exchange guarantees, not an unsupported claim about the logic of every P1 workload.

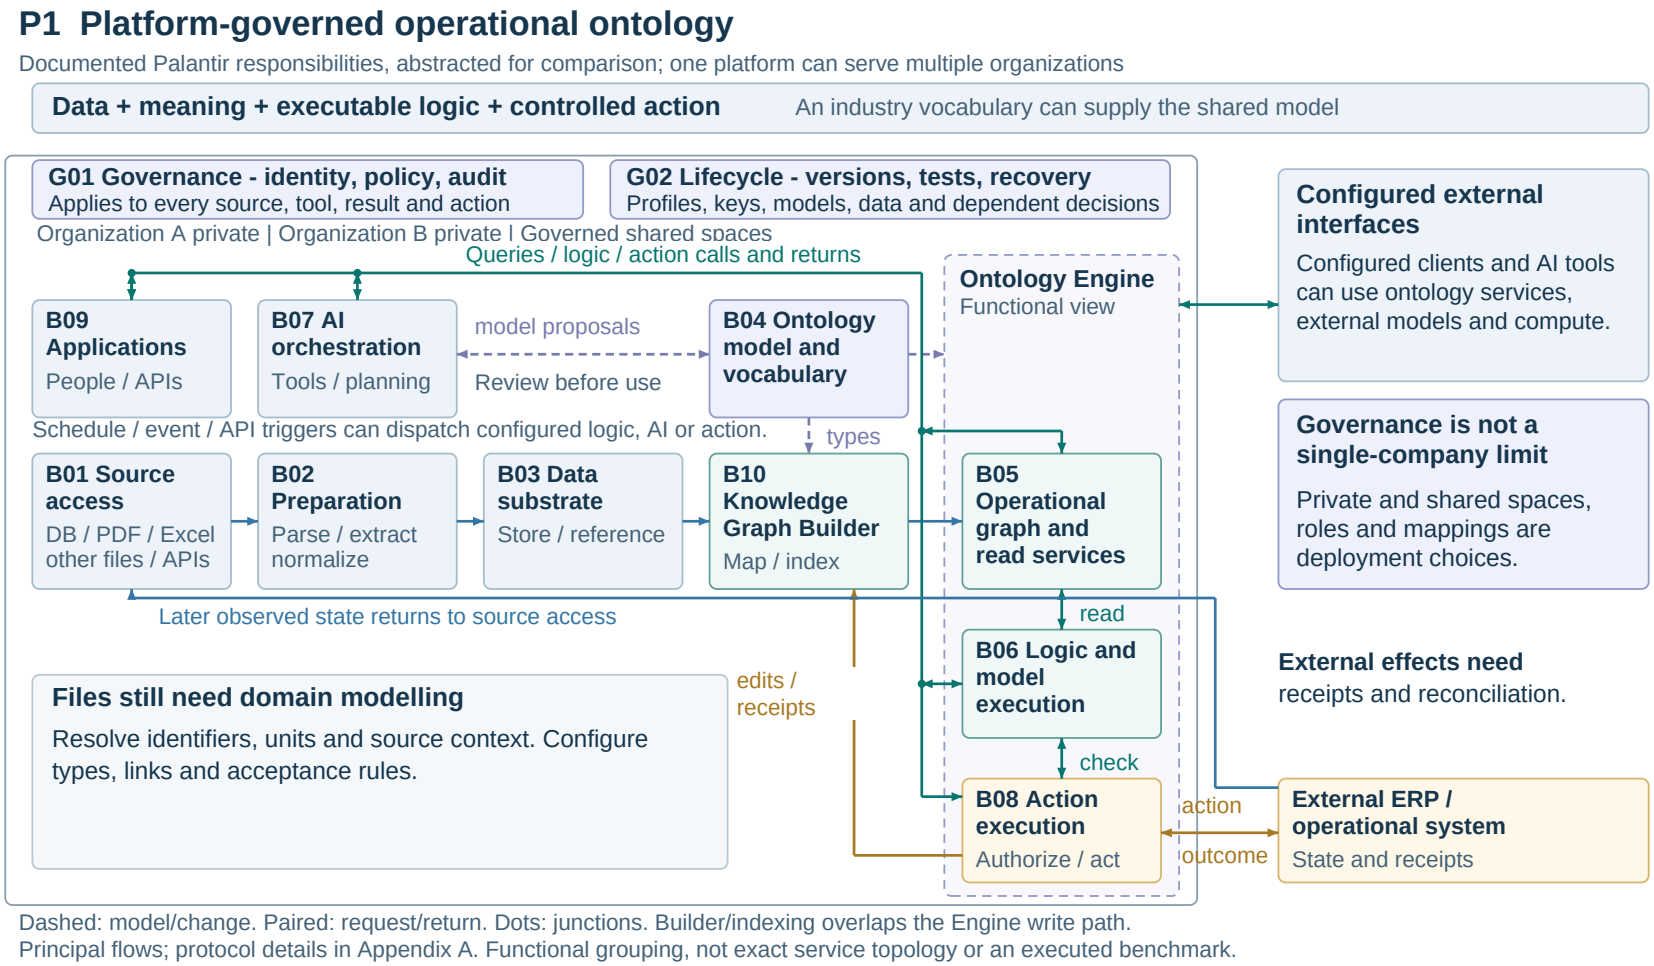


Figure 2: P1 integrates the twelve common responsibilities within a logical platform governance scope. The illustration permits private and shared organizational scopes and configured external services; it does not impose a single-company or single-database limit. The Knowledge Graph Builder is an author-defined functional grouping that overlaps the Ontology Engine write path. Principal data, configuration, request and action paths are shown; Appendix A records the detailed flow groups. Author synthesis from the documented P1 baseline.

5.3 AI uses governed tools and actions

Applications and AI orchestration use separate read, computation and action interfaces. An agent may search objects, traverse links, invoke a function and use the result to choose its next permitted tool. It need not load the whole graph into its prompt. Palantir documents ontology-aware tools and retrieval workflows, including configured confirmation behavior. Querying, numerical calculation, prediction, LLM explanation and formal inference nevertheless provide different kinds of justification. [29] [30] [31]

An authorized instance change follows B08 $\rightarrow$ B10 $\rightarrow$ B05: an action produces an edit, indexing incorporates it and later reads reflect the appropriate state. External operations follow another path to ERP or other operational systems. Palantir's documented atomicity for Ontology edits excludes external side effects. A lost ERP response therefore requires reconciliation rather than an assumption that the action either failed or completed. [32]

Engineering changes have a different authority boundary. AI FDE can propose changes through its documented branching and review workflow. That design-time path must not be confused with an ordinary instance write. Scheduled, event-driven and API-triggered automation can dispatch configured work without an LLM. Evaluation, versioning, promotion and recovery belong in the lifecycle architecture. [33] [29] [34]

5.4 Multi company operation and the portability question

Palantir explicitly supports multi-organization ecosystems with private and shared spaces; a common ontology can be industry-standard or participant-defined. It also documents APIs, imports, exports and external integration. A fair counter-model cannot rest on claims that Palantir only serves one company, cannot use consensus vocabulary or categorically blocks interoperability. [35] [18] [36] [37]

The stronger question is whether the required operational contract survives movement between independently implemented environments. An exported model may not preserve every policy, application or action behavior. This is a test obligation, not a presumption of failure. The Foundry connector's documented non-transfer of security markings between instances illustrates one specifically scoped migration condition requiring destination controls; it is not a universal finding about all Palantir sharing. [38]

P1's strengths include coordinated engineering and operational breadth. Its strategic risk, where present in a deployment, is that accumulated meaning, controls and applications may become costly to recreate elsewhere. The magnitude depends on actual configuration, contracts and available migration mechanisms. The counter-model must compete with P1's usefulness while making these dependencies inspectable and contestable.

6 Pattern P2 Participant controlled industrial cooperation

6.1 Preserve local capabilities and change the cooperation boundary

P2 retains all twelve functional roles within each participant. A supplier, laboratory and manufacturer may choose different graph engines, operational platforms and AI tools. Their cooperation depends on agreed external behavior rather than a requirement to operate the same internal stack. Palantir can supply one local implementation. The counter-model targets dependence on a common governing runtime, not the exclusion of competent commercial suppliers.

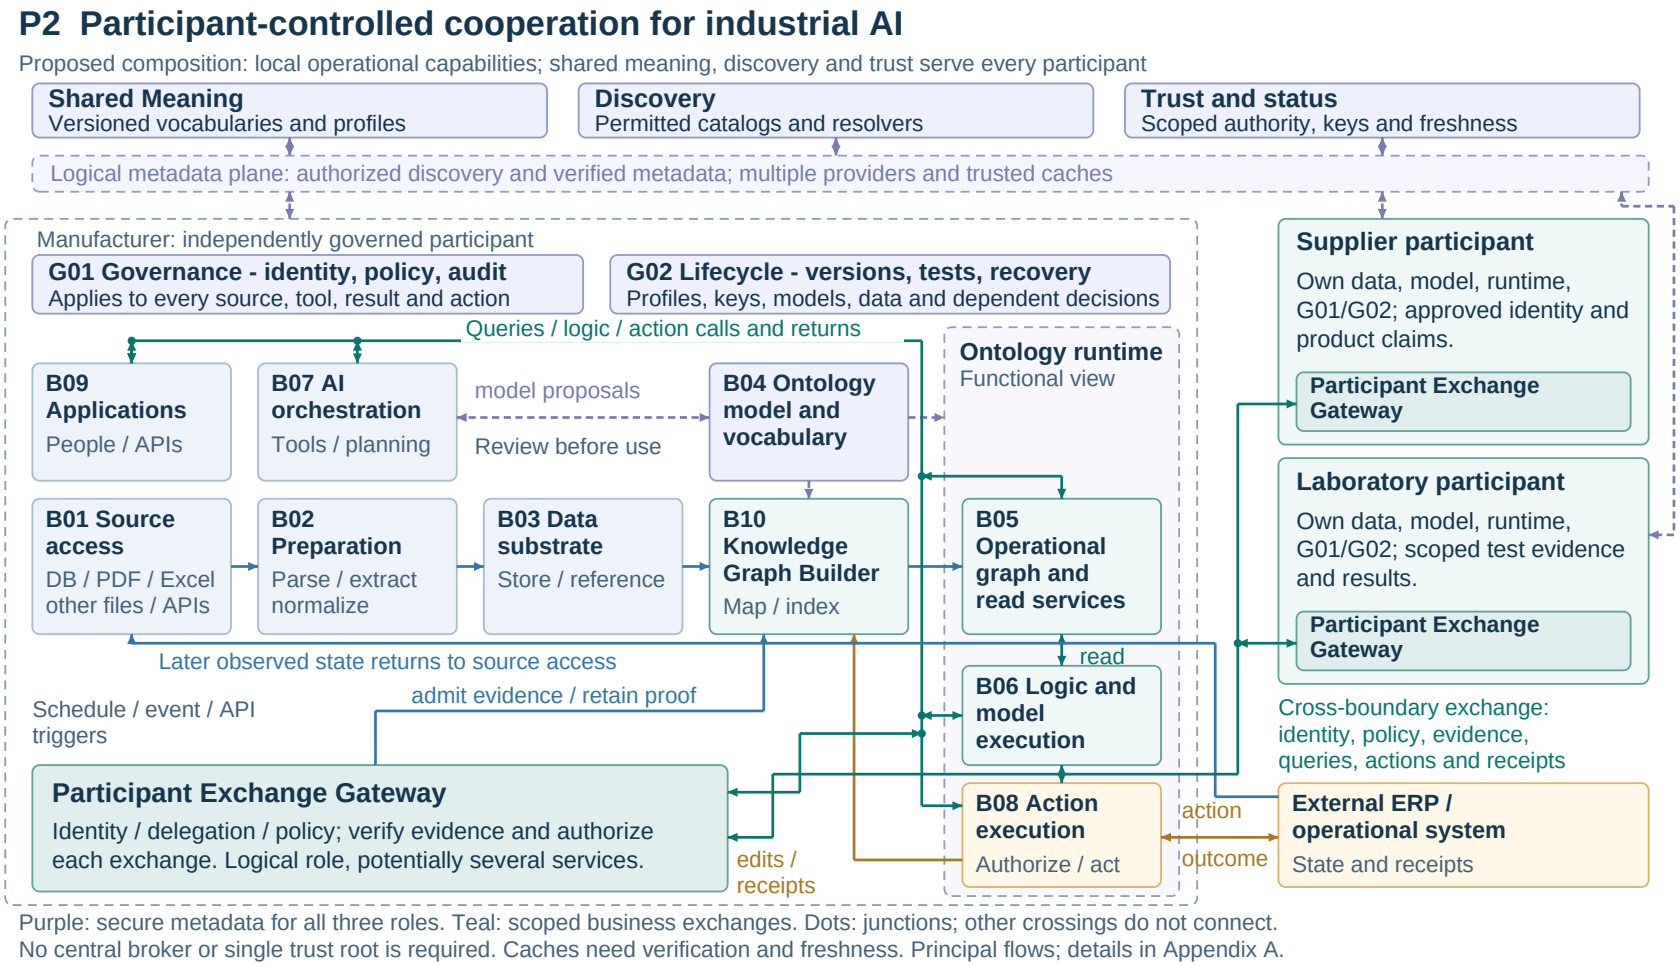


Figure 3: P2 repeats local model, builder and runtime responsibilities across independently governed participants. The manufacturer, supplier and laboratory each connect to all three shared roles: Shared Meaning, Discovery and Trust. The purple metadata connections are distinct from scoped business exchanges through Participant Exchange Gateways. A dataspace connector built with Eclipse Dataspace Components (EDC), complemented by identity and evidence services, is one gateway implementation example. Shared roles may have several operators; they imply neither a mandatory central broker nor unrestricted access. This is a proposed composition, not an executed integration. [39]

Figure 3 expands one participant and shows its peers through their exchange boundaries. Each peer retains its own model, builder, runtime, applications, governance and lifecycle. An agreed contract defines what crosses the boundary, who can assert or request it, what checks are necessary and what a failure means. This is the architectural realization proposed for the industrial meta-platform: reuse across platforms without transferring every local decision to one intermediary.

The **Participant Exchange Gateway** is the logical boundary for identity presentation, delegation, policy evaluation and scoped exchange. A **dataspace connector**, for example one built with Eclipse Dataspace Components (EDC), is an implementation option for its data-sharing functions. The broader gateway role also requires composed identity, evidence-admission, agent-authority and action-recovery services. EDC's connector control and data planes, Identity Hub and Federated Catalog have distinct responsibilities; installing a connector alone does not establish the complete P2 contract. [39]

A gateway may be managed by a service provider. Participant control means effective authority over the relevant policy, disclosure, key custody and exit arrangements, with appropriate contractual and technical guarantees. It does not require every participant to operate every server. The IDSA Rulebook's account of agency and autonomy provides a relevant governance reference, including the role and costs of intermediaries. [40]

6.2 Shared meaning discovery and trust

Shared Meaning publishes approved vocabularies, contexts, constraints and mappings. A semantic hub, ontology hub or vocabulary hub can perform this distribution role; the vocabulary is the content and the hub is a service. IDSA describes semantic provision through multiple hubs or references to existing resources. Shared Meaning does not impose one global ontology or confer truth on every statement using its terms. [41]

Discovery supplies permitted catalog and resolver metadata; DCAT is a relevant catalog standard. **Trust services** support accepted issuer scope, identities, keys and status under an identified authority. Catalog availability, shared vocabulary and accepted identity answer different questions. Their combination does not automatically grant access, establish a legal mandate or prove a product claim. [42] [43]

All participants need access to the shared roles relevant to their task. A supplier needs the approved product vocabulary, permitted discovery metadata and accepted trust information; a laboratory needs the corresponding measurement definitions, authority scopes and verification material. The manufacturer has the same responsibilities for its decisions. Each may also publish metadata where authorized. Figure 3 therefore connects every participant to a common logical metadata plane covering all three roles. It does not depict a service used only by the expanded manufacturer.

This is the metaplatform relationship: participants reuse capabilities across otherwise independent platforms. [7] In this design, maintained meaning, discovery and trust form a common logical plane of interfaces and governance responsibilities, not one server, universal vocabulary or compulsory broker. Different participants may use different authorized views and operators. Business evidence continues to flow through scoped participant exchanges. Metadata itself requires authorized access and verification against accepted publishers, including freshness rules for trusted local caches. No line grants access to every catalog, claim or trust record.

These services can themselves become concentration points. P2 therefore proposes portable releases, explicit governance, continuity arrangements and alternative operators where the task permits them. A community may legitimately maintain one agreed sector profile while several

operators distribute verified copies. Decentralized hosting and shared semantic authority are compatible when the authority and its limits are declared.

6.3 Admit evidence before using it for assurance

A knowledge graph relates entities and assertions. An evidence graph adds source attribution, verification state and relevant times. An assurance graph connects accepted evidence and assumptions to a bounded claim. These are working definitions; OMG's Structured Assurance Case Metamodel supplies a relevant formal reference for arguments, evidence and terminology, but this design does not claim SACM conformance. [44]

Admission checks the exact protected payload, issuer authority for the asserted scope, subject binding, validity and status, semantic-profile compatibility and fitness for the intended decision. PROV-O can express derivation and provenance; Verifiable Credentials and selected securing mechanisms support attributable assertions. None establishes the truth of an inaccurate measurement or the competence of an unaccepted issuer by itself. [45] [43] [46]

The original assertion and normalized graph view remain distinct. A value extracted from an unsigned PDF is still an extracted statement. Signing the extraction identifies the attesting party; it does not retroactively authenticate the original document's issuer. Likewise, transforming a signed payload into graph triples does not automatically extend its proof to arbitrary subsequent joins. Figure 5 shows the proposed admission and assurance relationship after the comparative analysis.

6.4 Reason over the authorized graph and act under local authority

The full industrial graph is a network of potentially connected knowledge. A particular computation sees the authorized, relevant and available portion. P2 permits disclosed evidence packages, remote queries, approved compute-to-data and hybrid snapshots. SPARQL federation supplies a query mechanism but does not supply all surrounding trust, confidentiality and completeness guarantees. A failed or denied request cannot be silently converted into evidence that a fact is false. [47]

B07 orchestrates permitted tools across participants. B05/B06 return results with source, snapshot and verification context. A transport-authenticated response may be useful for exploration, yet insufficient where a release decision requires a protected issuer assertion. Persisted, reusable query or computation results must pass the declared admission rules and enter the governed B08 $\rightarrow$ B10 $\rightarrow$ B05 path. Disclosed evidence packages can enter through the gateway and builder under their admission contract. The design does not require every transient query answer to carry a digital signature.

B08 retains local authority over consequential changes. Read permission for a laboratory summary does not authorize a manufacturer's agent to modify that laboratory's record. An action request binds the purpose, subject, operation, delegation and relevant evidence and policy versions. The receiving participant checks current preconditions. A receipt records the attempted or completed operation; subsequent observations establish external state when needed. Idempotency and reconciliation are requirements, not a promise of a global distributed transaction.

Status or profile changes can invalidate the current basis of a decision. Status-list mechanisms support revocation or suspension checks, but deployments must specify freshness, caching and outage behavior. The past decision remains part of attributable history where retention is permitted, while current acceptance and dependent workflows may require reassessment. Historical verification does not imply permanent authority. [48]

6.5 Policy interoperability and the secure semantic supply chain

ODRL provides a standard model of permissions, prohibitions, duties and constraints. A selected profile must bind those terms to policy decisions and enforcement at source, query, result, tool and action boundaries. EDC is a reference implementation route for an ODRL-based policy engine, but its behavior depends on scopes and registered functions; unbound rules can be filtered. P2 therefore requires an explicit supported-term gate before treating mandatory policy as enforceable. [49] [50]

Reasoning across sources also requires policy composition. A derived result may combine evidence subject to incompatible disclosure or retention conditions. The architecture must define conflict handling and defer operations whose obligations cannot jointly be satisfied. Policy expression and negotiation do not themselves prevent every downstream misuse or inference leak. Enforcement, monitoring, accountability and legal arrangements remain part of the design.

Semantic resources are security-sensitive dependencies. A substituted unit definition, poisoned context, stale issuer list or misleading catalog entry can alter a decision without changing its main source document. P2 applies resource-specific access and verification consistent with zero-trust principles. Publisher authority, immutable version identifiers, protected releases, dependency manifests and rollback/freeze detection are proposed controls. TUF offers relevant update-security mechanisms; applying them to semantic releases is our adaptation, not a standardized TUF semantic profile. [51] [52]

A cache improves resilience only within an explicit validity and freshness policy. Refetching an old status object does not make the original assertion current. DLT may support a selected shared record, but ledger inclusion does not alone establish issuer competence, acceptable status or complete knowledge. A verifier needs the approved authority and checkpoint rules for its particular task. Neither caching nor a ledger removes the need to stop when sufficient current evidence is unavailable. [53] [54]

7 Comparative analysis of control and dependence

7.1 A counter model with the same operational obligations

Figure 4 contrasts governance and dependence while retaining common operational responsibilities. P1 supplies documented integrated capabilities. P2 proposes cooperation through explicit boundary contracts and independent participant authority. The visual distinction concerns organization of control; it does not classify every centralized service as harmful or every decentralized component as safe.

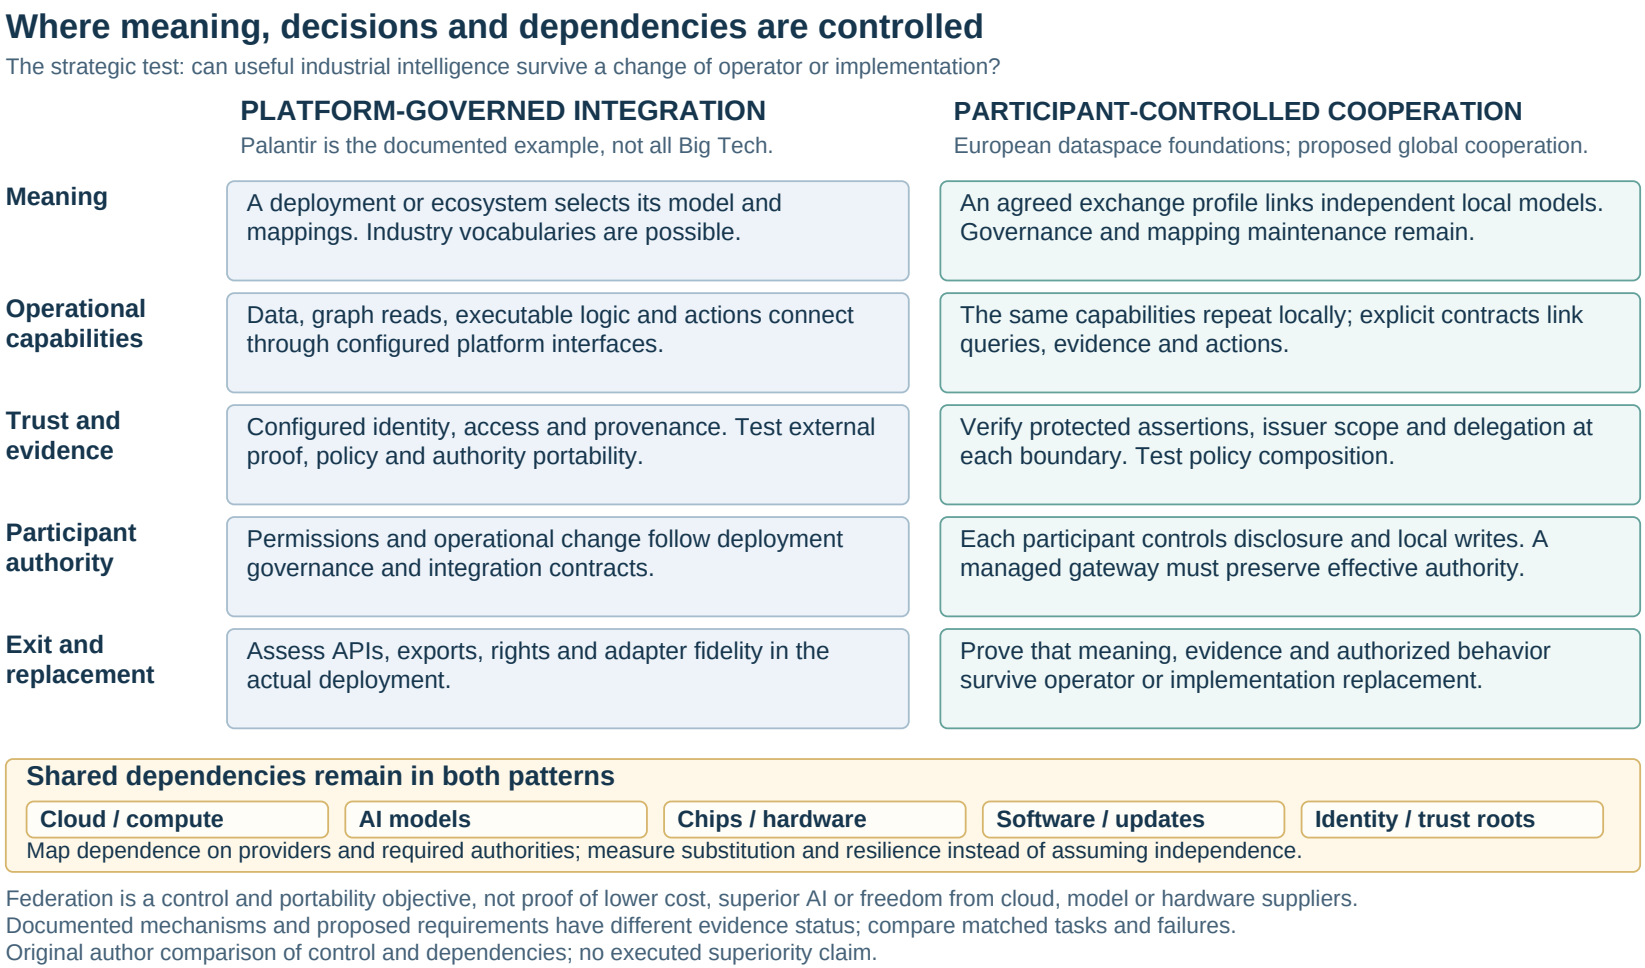


Figure 4: The strategic comparison is platform-governed integration versus participant-controlled cooperation across implementations. Both require meaning, evidence, policy, actions and lifecycle management, and both retain infrastructure and supply-chain dependencies. P1 is illustrated through documented Palantir capabilities, not a claim about all US Big Tech products. P2 is the proposed global cooperation model, building on European dataspaces and international contributions. Its differentiator is a requirement to demonstrate preserved behavior and credible substitution, not an assumed economic or security advantage.

The differences matter at four boundaries. A semantic boundary tests whether a recipient preserves the intended meaning. An evidence boundary tests whether an assertion remains attributable and acceptable outside its original environment. A policy boundary tests whether permissions and obligations retain their effect. An action boundary tests whether requests, approvals, outcomes and recovery remain understandable and authorized. Exporting data is necessary for some workflows but insufficient to demonstrate all four.

This analysis follows the meta-platform argument by making participant-controlled reuse a design objective. It extends that objective with operational conditions: a reusable statement must retain its interpretation, proof context, permitted use and relationship to action. A gateway that merely relays incompatible payloads does not create this capability. A common graph format without governed usage and validation leaves the hardest interoperability work unresolved.

7.2 Documented risks and changing provider behavior

The UK's 2025 cloud-market investigation found adverse effects on competition associated with concentration, switching and multi-cloud barriers, and Microsoft's licensing practices in the assessed markets. AWS and Microsoft are central providers in that analysis. These are cloud-market findings, not findings against Palantir or a common architecture attributed to all US technology firms. They demonstrate why an export API alone cannot establish practical choice: technical interfaces, software licensing and commercial arrangements also affect the ability to change providers. [55]

The CMA's 16 October 2026 update also reports provider improvements. Crucially for this paper, it identifies remaining costs of translating authorization logic across cloud identity and access-management languages and models. Identity federation and policy portability are therefore different achievements. This supports P2's requirement to preserve policy meaning across implementations while showing why the critique must be updated as suppliers improve. [56]

The EU Data Act addresses switching between data-processing services and establishes relevant interoperability obligations. Its legal categories and conditions matter; it should not be summarized as a universal guarantee that every application, model and behavior can migrate unchanged. Removing a contractual obstacle or a switching charge does not itself recreate a complex operational ontology elsewhere. Technical profiles and measured replacement tests remain necessary complements to legal rights. [57]

In the inspected announcement of 25 June 2026, the Commission expressed a preliminary position on cloud-gatekeeper designation; that source does not establish a final decision. This study does not infer a later outcome from the announcement. Regulatory scrutiny and documented provider improvements can coexist with remaining concentration and migration risks. A strategic architecture should respond to the mechanisms rather than depend on a fixed narrative of vendor misconduct. [58]

7.3 Geopolitical exposure and system level dependency

Industrial AI can become part of the infrastructure through which firms coordinate production, compliance and maintenance. Dependence becomes systemic when many organizations need the same indispensable service and cannot preserve essential functions through a disruption. In that setting, loss of access or support can have severe consequences across a supply network. Correlated outages and disruptive commercial or legal changes are plausible scenarios. This study does not estimate their probabilities or assert that a provider intends to cause them.

Jurisdiction is relevant but requires precise treatment. EDPB guidance on GDPR Article 48 addresses foreign public-authority requests and the conditions governing disclosure and transfer of personal data. It does not establish that every foreign request is valid, every dataset is personal data or every US-controlled service exposes all customer information. P2 does not automatically solve cross-border legal compliance either. The applicable entity, authority, data, control relationships and transfer rules must be examined. [59]

The resilience objective applies across regions: essential industrial operations should not depend on an untested assumption of indefinite access to one indispensable commercial stack. This is a strategic inference from the importance of the functions and the documented dependency mechanisms. Proportionate continuity and replacement capabilities can make international cooperation more durable, because participants can retain essential functions when commercial or political conditions change. The objective concerns effective authority and resilience, not insulation from foreign technology or an assumption that domestic ownership eliminates risk.

The Commission's sovereignty framework reinforces the multidimensional analysis. Hosting location, legal exposure, operational control and technological dependency are not interchangeable. A European federation can still depend on a non-European cloud, model API or update service. Conversely, a non-European vendor can supply a well-bounded component within a participant-controlled arrangement. Procurement must inspect the dependency chain and the authority exercised at each layer. [6]

7.4 Evidence and assurance as portable industrial capabilities

Figure 5 connects assertions to an explicit admission process, an attributed evidence graph, a bounded assurance argument and a separately authorized operation. This is the proposed bridge from semantic interoperability to useful industrial automation. It also defines a limit: an accessible explanation or a valid signature cannot independently establish functional safety, engineering correctness or legal acceptance.

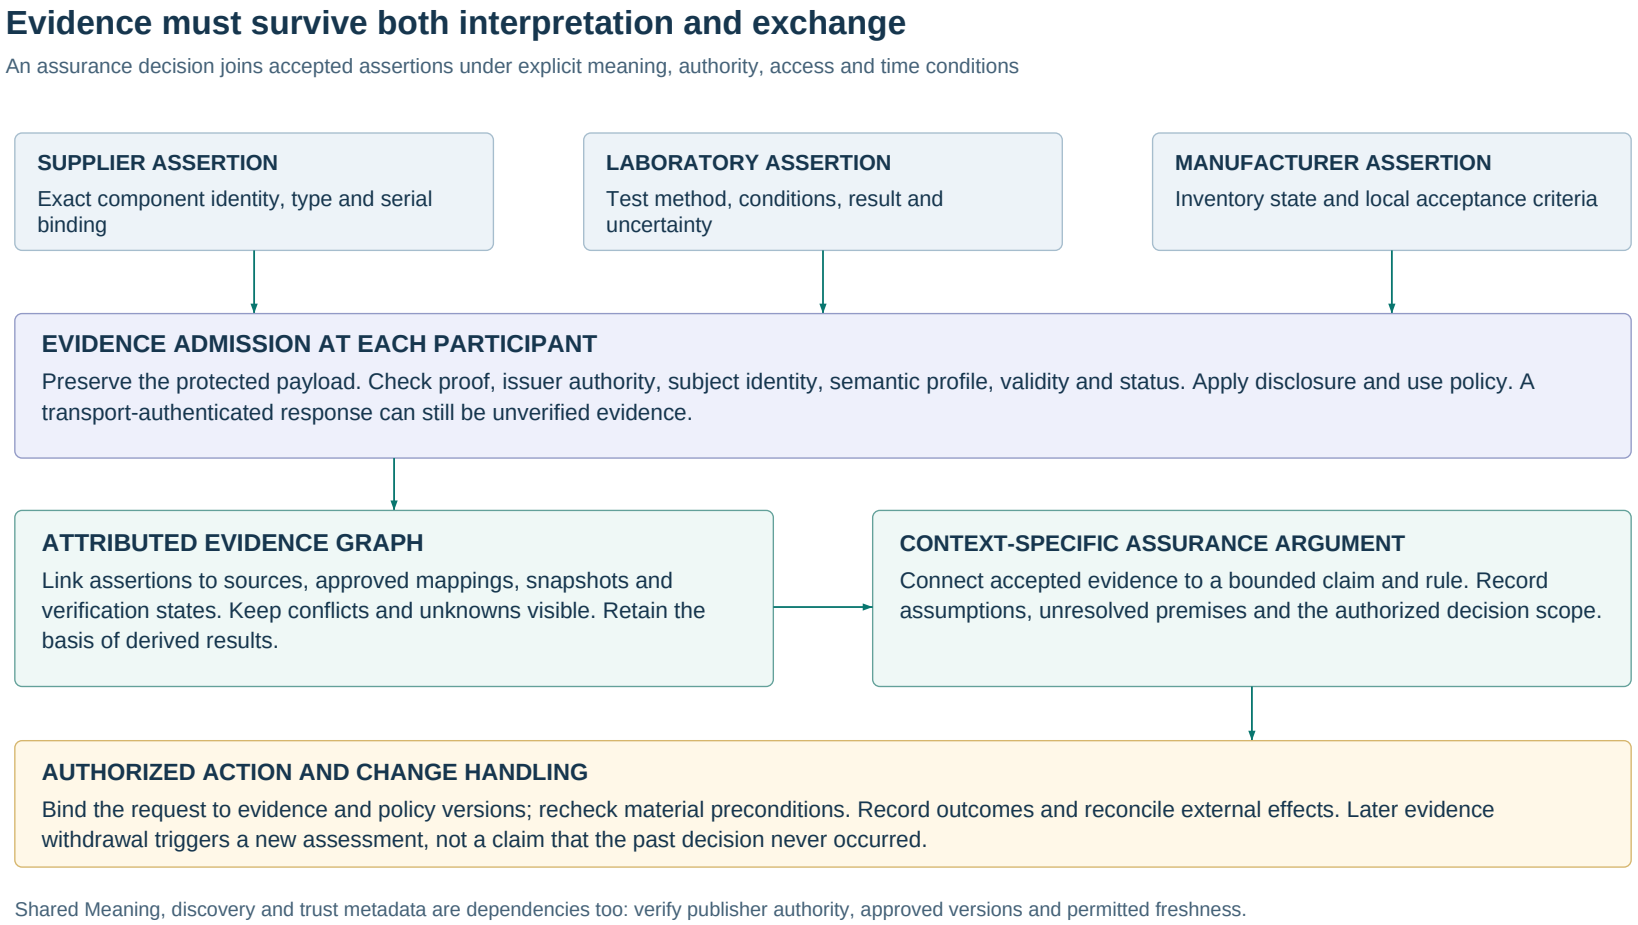


Figure 5: Evidence admission preserves the original assertion and checks its interpretation, issuer authority, status and permitted use before it supports a bounded argument. Authorized action remains a separate responsibility. The diagram is an analytical design, not a claim that a signature or graph proves a physical system safe. Status and profile changes may require reassessment of dependent decisions.

Physical AI makes these distinctions more consequential. A robot may combine component identity, software configuration, inspection evidence and delegated authority before requesting an operation. The same infrastructure can organize claims about physical products, virtual AI services and embodied systems, but each needs a domain-specific acceptance argument. The author's earlier industrial AI data-plane paper motivates this continuity; it is conceptual prior work, not independent experimental evidence for the present architecture. [4]

8 A humanoid robot example linking evidence to scoped decisions

The proposed humanoid H-01 illustrates the architecture through a non-medical industrial picking task. It is fictional: no physical test, issued certificate or cryptographic verification is reported. Figure 6 connects four independently governed participants: component supplier, laboratory, robot manufacturer and operator. The example draws on Spherity's PACE and robot-passport research, while retaining their distinction between evidence management and physical assurance. The existing MiPA viewer is a concept demonstration, not a verified humanoid deployment or manufacturer endorsement. [60] [61]

Three terms distinguish what the graph represents. **Evidence** includes the original report TR42, its method, provenance and current status. An **assertion**, TA42, states a bounded test result about configuration R3. **Assurance** is the structured argument that accepted premises support release claim RC3 for task TP-PICK under approval rule AR1. The report, assertion and conclusion remain separate objects. Firmware identity, the software bill of materials and installation records help determine applicability; they do not establish safety independently. An associated AI evidence-assistant service is a virtual service, distinct from the robot and its local safety controller.

A later status event RS9 withdraws TR42 from current reliance. Under the invented AR1 rule, the only accepted test support is then unavailable, so the release decision becomes HOLD pending reassessment. The historical report remains inspectable. Withdrawal neither proves the robot unsafe nor invalidates unrelated records. A replacement report requires fresh identity, scope, status and applicability checks. The same dependency principle can extend to skills and memory views, where permission to reuse experience remains distinct from permission to execute a physical operation. [62]

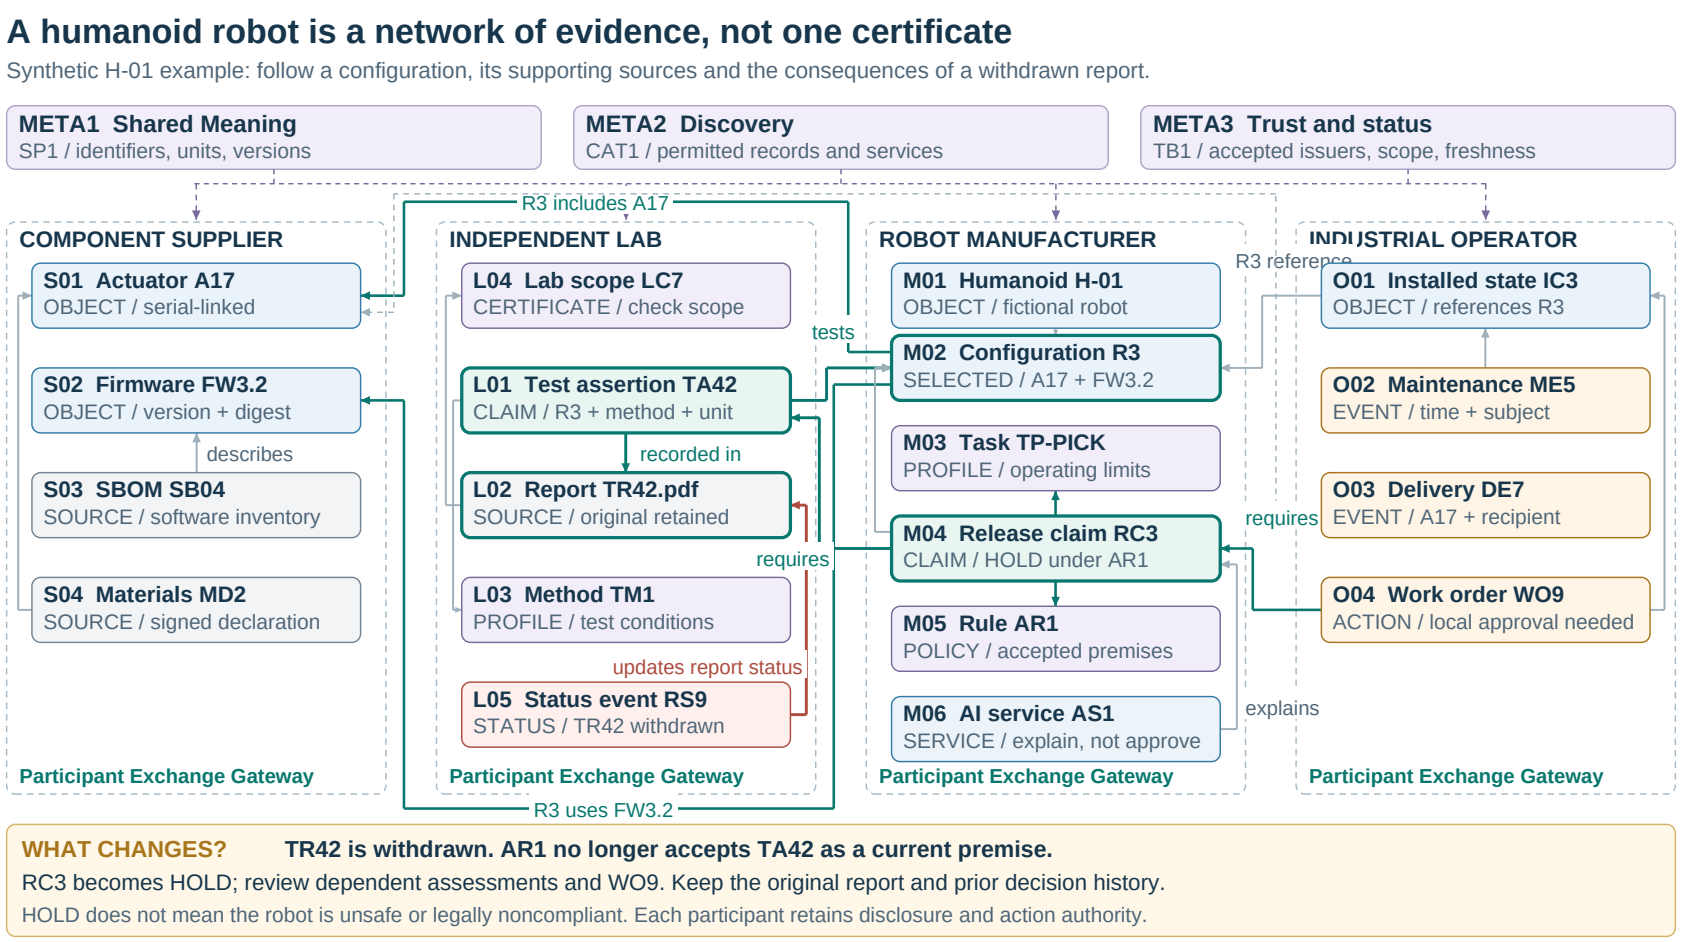


Figure 6: A synthetic humanoid evidence graph connects configuration R3 to components, a test assertion, its original report, scoped authority and operational events across four independently governed participants. Shared semantic, discovery and trust roles supply authorised metadata to all participants; depicted cross-boundary relationships are logical references, not unrestricted data flows. The highlighted trace shows a withdrawal of report TR42: under the illustrative rule AR1, release claim RC3 becomes HOLD and dependent assessments require review. This preserves historical records and does not establish that the robot is unsafe or legally noncompliant. All identifiers, sources, credentials, events and outcomes are fictional; no signatures or engineering results have been validated.

Figure 7 shows seven assessment lenses: ESPR/DPP, functional safety, the EU AI Act, MDR, machinery, CRA and export controls. Shared evidence can inform several assessments, but their scope tests, responsible parties and approval routes remain distinct. An ESPR passport duty depends on applicable product measures. Functional safety requires system and application assessment. [63] [64] AI classification and the current machinery-sector route require their own analysis. [65] [66] [67] A medical intended-purpose variant is a separate branch; the CRA excludes products covered by MDR. [68] [69] Export screening considers classification, destination, end user and end use. None follows automatically from the word humanoid. [70]

Each company retains its records and authorization boundary. A requester receives a task-scoped evidence view with provenance, use conditions and unresolved gaps; Shared Meaning, Discovery and Trust support all participants without requiring pooled files. AI may retrieve dependencies, explain affected claims and prepare an assessment dossier. Final acceptance and consequential action remain with authorized actors. Such benefits are conditional on correct mappings, evidence quality and measured performance; industrial AI does not universally require a knowledge graph. Candidate benefits to test are compliance automation, lower assessment costs, shorter time to market, better risk and cybersecurity management, more resilient recovery, and faster product and service innovation. The illustrated reuse is a proposed evaluation subject, not a validated compliance result.

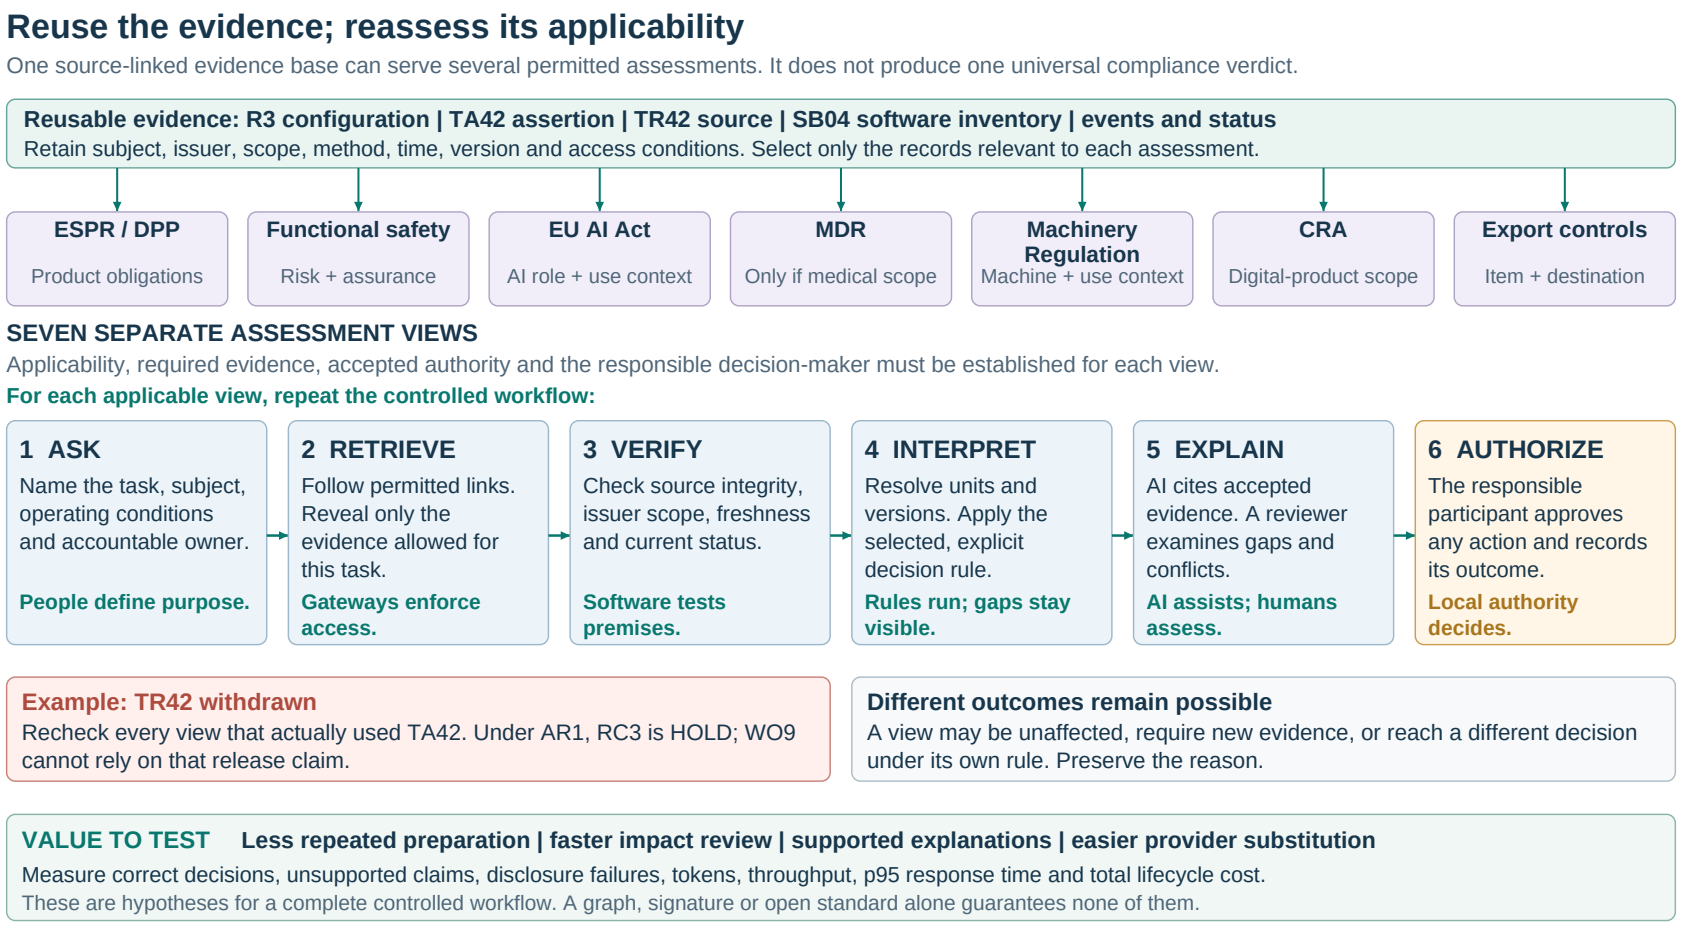


Figure 7: Selected evidence can be reused across seven separate assessment views: ESPR/DPP, functional safety, the EU AI Act, MDR where medical scope applies, the Machinery Regulation, CRA and export controls. Each view requires its own applicability analysis, accepted evidence, authority and decision rule; the arrows do not assert that every record is relevant to every regime. The proposed workflow separates permitted retrieval, verification, interpretation, AI-assisted explanation and locally authorised action. A withdrawn source triggers review only where it supports a dependency. Reductions in repeated work, response time or cost and improvements in explanation or portability are hypotheses to measure, not observed outcomes or automatic compliance.

9 An industrial example and a falsifiable comparison

The synthetic case concerns one identified valve, supplied in Japan, tested by a laboratory in the Republic of Korea and assessed by a manufacturer in Germany for a demonstration water loop. These locations illustrate organizational and jurisdictional separation. They confer no real accreditation, statutory recognition or transfer permission. The case contains no actual company data or cryptographic credentials.

Under the invented DEMO-LT-1 procedure, collecting 2.40 mL over 600 seconds yields 0.24 mL/min. Adding the stipulated uncertainty of 0.02 mL/min gives a guarded result of 0.26 mL/min against the scenario limit of 0.50 mL/min. These values are an explicit fixture, not an engineering standard or reported physical test. The manufacturer must also possess acceptable identity, authority, status and policy evidence. A favorable measurement alone cannot authorize the workflow.

Four outcomes keep uncertainty visible. RELEASE_ELIGIBLE permits a separately approved workflow request. REJECT_ENGINEERING means valid and interpretable evidence fails the numerical rule. HOLD means required premises are missing, denied, invalid, conflicting or unsupported. NO_NEW_DECISION rejects a consumed-request replay. The only permitted side effect is a controlled inventory transition; physical actuation and statutory certification are outside the experiment.

The fifteen planned cases cover healthy behavior, absent evidence, denied access, contradictory claims, revocation, wrong issuer authority, unsupported mandatory policy, stale status during outage, unsupported remote conclusions, uncertain ERP outcomes, replay, provider replacement, payload tampering, an unapproved semantic profile and an engineering failure. Each has a held-out oracle and declared disclosure and side-effect expectations. The replacement case binds a fresh request to the replacement service while preserving the laboratory's accepted authority.

Three functional repetitions produce forty-five planned attempts per pattern. The eleven blocking cases account for thirty-three of those attempts. No majority vote can hide a failed authorization, confidentiality or side-effect condition. The protocol records supported native behavior, required adapters, unsupported configurations and unassessed capabilities separately. Appendix C provides the execution boundaries; the detailed fixtures remain machine-readable.

Only after deterministic correctness should the study add AI conditions with fixed model versions, prompts, permitted tools and budgets. Compare retrieval assistance, local graph assistance and cross-participant graph assistance under the same admission and action gates. Repeat stochastic runs and report variation. A more persuasive explanation must not count as a better result if it invents authority, hides missing evidence or attempts an unauthorized action.

The resulting research propositions are falsifiable. Agreed profiles should preserve decision-relevant meaning across implementations. Provider substitution should preserve accepted evidence, policy effect and decision history with a measurable integration cost. Cross-participant AI should improve a declared task relative to its controls without increasing prohibited disclosures or actions. Failure on any of these dimensions weakens the counter-model's practical claim, even if all components use open standards.

10 Existing foundations and unresolved implementation work

10.1 European coordination with international participation

Europe has a substantial basis for this strategy, but maturity varies. The Commission's Common European Data Spaces overview covers fourteen sectors or domains; that is not evidence of fourteen universally interoperable production systems. Simpl's inspected development update describes Gdynia as available while listing Hamburg as planned. An elapsed target date is not proof of delivery. [2] [71]

Factory-X reports completion of consortium work on 30 June 2026. Catena-X reported 1,069 actively exchanging OEMs and suppliers on 21 September 2026, distinct from its 195 association members. These are attributed evidence of industrial activity. They do not establish arbitrary platform substitution, independently audited economic gains or European leadership on every architectural dimension. [72] [73]

IPCEI-AI supplies a policy and investment direction. On 16 September 2026, the Commission welcomed design work by nineteen Member States and reported planned pre-notification by eleven. Ireland's 30 October announcement adds political support and a statement of national funding allocation. Those sources do not establish Commission state-aid approval, deployed federated AI services or a mandate to use RDF or OWL. Programme association does not qualify existing dataspace components for demanding AI workloads. Investment should support necessary redesign and independent workload testing, with reuse justified by results. [74] [75]

The international contributions are substantive. IPA and Catena-X reported a battery carbon-footprint proof of concept on 31 March 2025 using a replicated Japanese environment and an intermediate translation layer. China's national data authority reports sixty-three trusted-data-space pilots in its development report dated May 2026. The first is bounded exchange evidence; the second is evidence of a national programme, not proof of identical semantics or production compatibility with P2. [76] [77]

The inspected UNTP specification remains described as work in progress for pre-production pilots, with a future 1.0 expected. W3C standards and open implementation communities contribute across regions. The European leadership proposed here is the ability to organize procurement, governance and industrial implementation around interoperable participant control. It is compatible with international contribution and should be evaluated through outcomes rather than claimed from the nationality of standards contributors. [78] [43]

10.2 Fast-track modernisation for AI-first operations

Modernisation must address the complete operational path: discovery, authorization, evidence retrieval, model execution and locally approved action. High throughput, low latency and large data volumes describe different requirements; each needs a declared workload and acceptance threshold. A document-exchange pilot cannot establish readiness for concurrent agents, frequent updates or time-sensitive industrial decisions. Conversely, a demanding AI workload need not route every byte through a graph or connector. Bulk files, tensors and raw telemetry can remain in appropriate object stores, streaming systems and time-series databases, while graphs connect identifiers, semantic context, derived claims, provenance and evidence-window references. Those data paths remain subject to agreed access, integrity and evidence controls. This follows the prior evidence-graphs proposal, whose industrial performance claims remain to be tested. [4]

We propose readiness gates for sustained and burst capacity, end-to-end median and tail latency, update and revocation propagation, recovery under failure, and total cost per completed authorized task. Tests should state payload mix, graph size, update rate and concurrent demand, including cold caches and unavailable trust services. Reuse of verified metadata or scoped authorization requires explicit validity conditions and reassessment when relevant status or policies change. Improved speed cannot compensate for prohibited disclosure or unauthorized action. Meeting performance and trust requirements together is the objective; it does not imply cost-free verification or the absence of engineering trade-offs.

These gates extend the research agenda beyond the small synthetic fixture in this paper; no such load test has been executed. International partners should agree profiles and representative workloads, test independently and publish limits as well as successes. Existing components are candidates for reuse, redesign or replacement. A successful pilot supports a bounded readiness claim for the conditions and properties actually tested, not universal readiness for IPCEI-AI. Sustained work on runtime design, semantic stewardship, security and operations is required to turn the first step into dependable infrastructure.

10.3 Standards do not compose themselves

Candidate participant stacks use Jena and RDF4J with separately implemented application adapters. They initially share EDC and credential-verification components. That design can test graph and application diversity; it cannot demonstrate independent implementations of every connector or cryptographic protocol. Exact candidate versions, runtime requirements and discrepancies belong in the supplementary component map. [79] [80] [81]

A concrete gap concerns credential securing formats. Digital Bazaar's VC library records VC 2.0 support but excludes JWT serialization in its README. UNTP's versioned 0.7.0 credential profile requires JOSE enveloping. A Data Integrity-oriented library is therefore not automatically a complete UNTP verifier. The rolling UNTP material includes proposed JOSE/COSE requirements that must not silently replace a frozen profile. [82] [83]

Dataspace Protocol and Decentralized Claims Protocol define useful exchange and participant-claim mechanisms, but trusted-issuer governance, compatible profiles and application-specific evidence admission still need composition. A protocol version, connector release and credential library are candidates until the required combination has been tested. The current supplement records integration gaps rather than a deployment-ready bill of materials. [84] [85]

The same discipline applies to public-authority operations. Submission, registration, inspection, correction and approval are distinct actions with different mandates. An identity credential or ODRL permission does not supply the legal authority for every operation. A B2G profile must join accepted institutional authority, applicable rules and tested access control. The architecture supports that work; it does not confer mutual recognition across jurisdictions by technical fiat.

11 Geopolitical context and digital diplomacy

11.1 A global counter model based on reciprocal agency

Industrial supply chains need cooperation across boundaries that no single company or jurisdiction controls. A test assertion may originate in Asia, enter a European assurance process and be interpreted by software developed in the United States. Requiring every contributor to use the same operational platform can simplify integration while placing considerable authority in its operator. Requiring all evidence to remain inside disconnected regional systems can make

legitimate trade and coordinated risk management unnecessarily costly. The proposed alternative makes selected relationships interoperable while leaving consequential decisions with their accountable participants.

We use **global digital sovereignty** to mean the effective capacity of participants in different jurisdictions to govern their digital relationships and cooperate under applicable rules. It is a design and governance objective, not a new legal status. Public authorities, firms and individuals have different powers and rights. A business policy cannot override a person's legal rights, and a national policy does not make every requested data use legitimate. A multipolar digital economy needs institutions through which these interests can be expressed, contested and enforced.

The counter-model promotes fair competition by making useful interfaces implementable by several providers. It supports resilience through bounded dependencies, attributable evidence and tested replacement. Its diplomatic value is reciprocal agency: each side can assess a partner's claims without requiring wholesale adoption of that partner's platform or regulatory system. Those mechanisms are proposed benefits to evaluate. Distributed infrastructure can still be captured by powerful intermediaries, and its costs must be assessed alongside the coordination benefits of integrated platforms.

11.2 European leadership through shared implementation

The EU's International Digital Strategy of 5 June 2025 proposes connecting international partnerships through technical exchanges, common projects and interoperability pilots. It also supports common specifications and open standards for digital public infrastructure. The Council's conclusions of 20 November 2025 reinforce partnerships, secure infrastructure and resilient supply chains. These are documented policy foundations for collaboration; neither document adopts this paper's architecture. [5] [86]

The Commission's Data Union explanation likewise connects sovereignty with trusted cross-border exchange and describes data labs intended to link dataspaces with AI ecosystems. This connects Europe's industrial data agenda to the present architectural question: how can information become useful to AI while its meaning, authority and permitted use survive organizational boundaries? The policy direction motivates the work without establishing that the required operational interoperability has been achieved. [87]

We propose participant-controlled industrial infrastructure as a core practical pillar of EU digital diplomacy. Europe's contribution would be to help finance and convene joint profiles, conformance work, implementation capacity and dispute mechanisms. Partners should help define requirements and govern releases rather than receive a finished European blueprint as a condition of participation. European leadership can then be evaluated by the breadth and usefulness of cooperation it enables, including competitive participation by firms outside Europe.

This approach should extend beyond a coalition of large economies. Smaller suppliers and lower-capacity administrations need affordable participation, documentation in usable languages, local implementation options and support for maintaining evidence. Participation costs and representation should be explicit procurement criteria. Otherwise, ostensibly open requirements can become a market-access burden or give dominant integrators greater control. This is a recommended governance condition, not a finding that current programmes already satisfy it.

11.3 UNECE traceability and transparent standards governance

The United Nations Economic Commission for Europe, or UNECE, provides a multilateral setting for trade facilitation and electronic business through UN/CEFACT. Its work offers a useful

bridge between regional industrial initiatives and wider supply-chain practice. UNECE Recommendation No. 46 addresses traceability and transparency in garment and footwear value chains. Its relevance is the organization of attributable supply-chain information and cooperation; the recommendation does not by itself prescribe the complete industrial-AI architecture proposed here. [88] [89]

Institutional status matters. In 2025, UN/CEFACT approved Recommendation No. 49 after deciding to remove references to UNTP and Annex I. The adoption must therefore not be described as UN/CEFACT approval of UNTP. Recommendations 46 and 49 provide policy guidance; they do not themselves create binding DPP obligations. This distinction supports a diplomatic principle: shared policy objectives, project specifications and formally approved instruments need identifiable governance and separate claims of authority. Compatibility should be established through transparent profiles and tests, without borrowing institutional legitimacy from an adjacent initiative. [90]

UNTP remains a relevant technical project for decentralized supply-chain evidence. Its inspected version register labels version 0.7.0 for pilots, while rolling material is work in progress and a future 1.0 is expected. Its guidance also distinguishes technical exchange from the legal and institutional conditions for cross-border acceptance. These are useful inputs to P2, with the limitations of their stated maturity. Alignment with UNECE traceability objectives and participation in UNTP work are complementary propositions; they are not interchangeable claims of formal endorsement. [91]

The wider initiative is not confined to Europe. Standards Australia and the International Code Council announced collaboration on a Universal Data Protocol for the built environment on 16 November 2024. The proposed UNTP extension and sustainability vocabulary connect international partners and Asia-Pacific standardization expertise with the same interoperability problem. An announced development project is evidence of cooperation and intent, not proof of deployed global recognition or conformity to a final UNTP release. [92]

11.4 Asia Pacific partnership and regional choice

Asia-Pacific encompasses different industrial structures, legal regimes and technical initiatives. A compatible cooperation model must accommodate that diversity through agreed boundary contracts. It should offer a route for regional systems to retain local governance while exchanging selected, attributable information. The objective is a bridge for lawful cross-border trade, with each participant able to evaluate purpose, authority and restrictions.

Japan provides concrete evidence at two levels. IPA and Catena-X reported a battery carbon-footprint interoperability proof of concept on 31 March 2025, using a replicated Japanese environment and an intermediate layer to reconcile differences. The result is bounded but architecturally instructive: cooperation did not require identical internal platforms. On 15 October 2025, IPA, the Data Society Alliance, the Robot Revolution and Industrial IoT Initiative, and the University of Tokyo announced coordination around Open Data Spaces specifications. That agreement supports an international interoperability agenda; it does not establish that all Japanese deployments use one architecture. [76] [93]

China's Trusted Data Spaces Development Action Plan for 2024–2028, issued on 21 November 2024, describes a national programme that includes cross-border cooperation and international standardization. It is evidence of strategic policy support, not merely an isolated industrial experiment. An official notice also reports that China's data-standardization committee, SAC/TC609, and the International Data Spaces Association signed a cooperation memorandum on 17 December 2025. This is a specific institutional bridge for standards work. IDSA

is an industry association; the memorandum is not an EU treaty or a finding of operational compatibility. [94] [95]

A Chinese AI cooperation plan published on 17 July 2026 additionally connects cross-border trusted dataspaces with high-quality datasets, international open-source cooperation and standards coordination. These stated aims provide a basis for identifying shared technical work, while leaving implementation outcomes to be demonstrated. They do not establish identical RDF/OWL semantics, credential profiles or authorization behavior across Chinese and European systems. [96]

The diplomatic distinction is between political alliance, standards cooperation and acceptance for a particular transaction. A Chinese institution can participate in an appropriate standards or research activity without being characterized as an EU ally or being granted access to all industrial evidence. Similarly, a European or US provider does not gain universal acceptance from its location. Scope, authority, applicable rules and tested behavior should determine the technical relationship. This makes cooperation selective and accountable while avoiding nationality as a substitute for architectural analysis.

11.5 US innovation and global web standards

US innovators are contributors to the counter-model, not simply suppliers against which Europe must defend itself. W3C's Verifiable Credentials Data Model 2.0 credits editors affiliated with Digital Bazaar, OpenLink and Block alongside wider expert participation. An official US small-business award record identifies Digital Bazaar in Virginia and a DHS award for credential and wallet research and proposed first-responder pilots. This supplies a concrete example of US innovation in shared infrastructure; the award does not establish completed results or imply US government endorsement of P2. [43] [97]

The W3C Verifiable Credentials Working Group's 2026 charter includes tentative DPP and business-wallet vocabulary work drawing on inputs including CIRPASS-2, WE BUILD and UNTP. Its DPP and Business Wallets Task Force provides a forum for use cases and requirements. This is a practical connection between European industrial and wallet initiatives and a global standards community. A charter, task-force discussion or use-case document has a different status from an adopted W3C Recommendation. [98] [99] [100]

The Recognized Entities Working Draft dated 6 September 2026 illustrates the next layer of the problem: representing recognition and authority relationships so a verifier can assess who is accepted to make which claims. Its contributors span organizations including Digital Bazaar, Fraunhofer IAO, TNO, Keio University and Spherity. It remains experimental and is expressly not presented as production-ready. Encoding a recognition relationship does not create a treaty, accreditation or statutory mandate; the relationship must refer to an accepted authority and scope. [101]

Established W3C standards already provide reusable parts of the design. RDF and JSON-LD support shared representation, Verifiable Credentials support attributable claims, and ODRL expresses permissions, prohibitions and duties. Their use can connect implementations developed in different regions, but a shared specification is only one layer of compatibility. For example, ODRL requires processing of a policy to stop when an asserted profile is unrecognized; an implementation must bind the selected vocabulary to actual policy evaluation and enforcement. This turns standards cooperation into a concrete obligation to test behavior. [15] [16] [43] [49]

11.6 Digital diplomacy through bounded cooperation agreements

Figure 8 summarizes the proposed institutional architecture. Participants retain local operations and authority; jointly governed interfaces make selected claims and requests usable across those boundaries. Shared semantic, discovery and trust services support the exchange, but do not become one global decision-maker. The figure describes a design for collaboration, not an existing treaty, alliance or completed interoperability network.

A bounded cooperation agreement should specify the business task, semantic and policy profiles, accepted issuer scopes, verification requirements and authority for consequential actions. It should identify permitted hosting and processing arrangements, applicable confidentiality and transfer requirements, correction and appeal procedures, and continuity arrangements. Each party needs a way to challenge a claim or refuse an operation when the necessary conditions are absent. These are proposed design requirements for digital diplomacy translated into industrial practice.

Evaluation should distinguish three outcomes. Technical interoperability means a selected exchange preserves required representation and behavior. Institutional acceptance means the recipient accepts a specified issuer and claim for a defined purpose. Legal recognition concerns the effects granted under applicable law or an agreement. Success at one level does not establish the others. Cooperation can begin with a narrow non-sensitive case, record unresolved issues and extend its scope only when the relevant conditions are established.

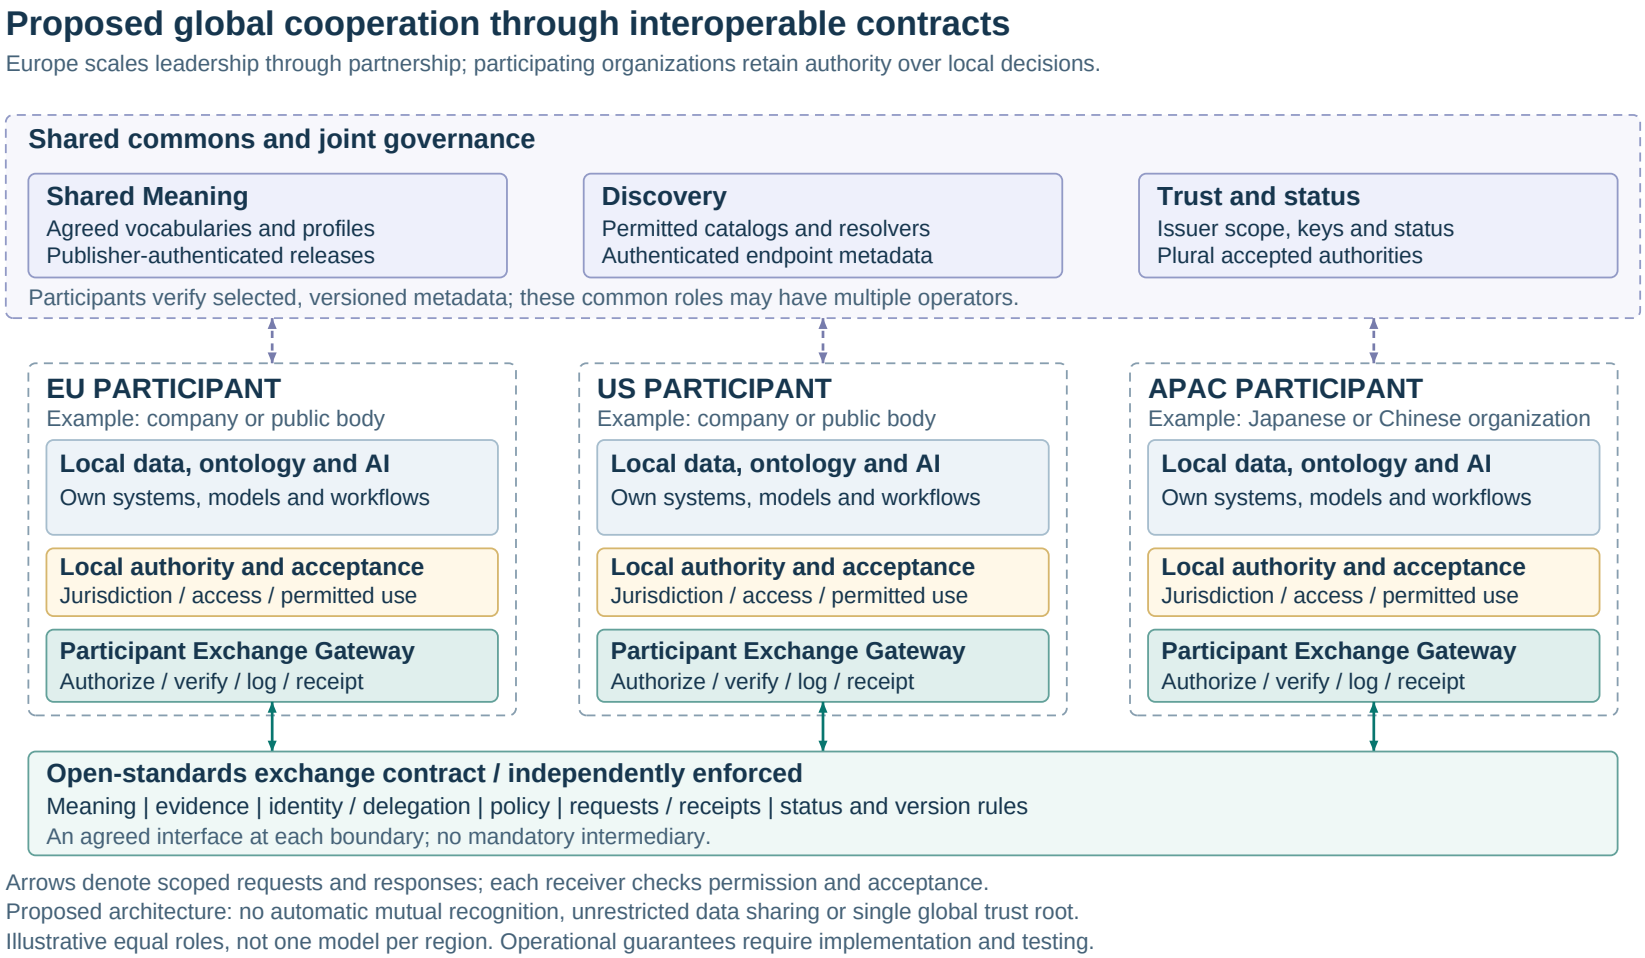


Figure 8: Proposed global cooperation through open exchange contracts and participant-controlled gateways. European dataspace experience, US innovation and Asia-Pacific initiatives can contribute to shared meaning, discovery and trust under joint governance. Regional participants retain local authority and applicable jurisdictional requirements. Technical interoperability, accepted issuer authority and legal recognition require separate evidence; arrows represent permitted exchanges and governance relationships, not unrestricted data movement or a completed international deployment.

The practical next step is a jointly governed experiment with participants from several jurisdictions and replaceable service providers. Use shared conformance artifacts and publish the supported profiles, mapping losses, policy decisions and failure cases. Measure onboarding burden and benefits for each participant type, including smaller firms. Such work would turn Europe's dataspace experience into a global collaborative offer whose value partners can verify and help shape.

12 An international industrial and procurement strategy

12.1 Fund cooperation as an operational capability

The strategic necessity is to retain agency over infrastructure that increasingly interprets industrial reality and initiates consequential workflows. Europe and its partners should make preserved meaning, evidence and behavior a procurement objective, alongside price, performance and security. Funding another catalog or graph database is insufficient if firms still cannot connect their operations without rebuilding the same bilateral integrations or surrendering practical choice to a new gatekeeper.

The proposed counter-model has three mutually dependent investments. Maintained semantic and policy profiles make agreement concrete. Conformance artifacts and reference adapters make implementations comparable. Sustainable governance determines who can change definitions, accept issuers, resolve disputes and operate shared services. These are ongoing industrial functions, not one-off standardization deliverables. Their financing and accountability need to survive beyond a pilot grant.

Public programmes can lower coordination costs by procuring complete, bounded use cases across several independent suppliers. Each procurement should specify required outcomes and portability properties without mandating one internal database or AI model. It should provide a route for capable incumbents, including Palantir, to participate through verifiable contracts. This keeps the counter-model focused on contestable cooperation rather than replacing a foreign platform with an equally indispensable domestic platform.

12.2 Require an exit that preserves useful behavior

A credible exit exercise replaces a selected service while preserving the required evidence, meaning, policy effect and decision record. It measures what changed, how long the work took, what information was lost and which functions could not be reproduced. Merely delivering a file export should not satisfy a requirement whose purpose is continuity of industrial operations. The assessment should distinguish an optional intermediary from a mandated institutional authority.

Procurement should identify several layers of exit. Data and model exports concern retained assets. Semantic exit concerns identifiers, definitions, mappings and versions. Evidence exit concerns original assertions, verification context and audit history. Policy exit concerns translation and equivalent enforcement. Operational exit concerns applications, actions and recovery. Infrastructure exit concerns the clouds, models, software updates and key services on which all the preceding layers depend. These are proposed assessment dimensions, not assurances that every layer can be made interchangeable at low cost.

The CMA's recent findings give policy portability particular urgency. A portable organizational identity does not automatically reproduce resource permissions or authorization logic. Procurement should require concrete tests of selected policy equivalence and unsupported terms, with

transparent disclosure of custom translation work. It should also recognize provider improvements when they reduce the measured burden. A counter-model grounded in outcomes remains useful even when incumbent platforms become more interoperable. [56]

12.3 Govern the commons without centralizing the market

Shared vocabularies and trust infrastructure need ownership, maintenance and liability arrangements. Open publication alone does not decide whose interpretation prevails during a dispute or who pays for a critical update. Governance should define representation, review, extension, version retirement and appeal. SMEs need usable services and proportionate integration work; forcing every participant to become a semantic-web specialist would undermine adoption.

The meta-platform contribution is to make cooperation reusable without requiring all value creation to occur inside one platform. Its industrial extension should be judged by successful relationships and retained choice: time to onboard, accepted claims reused, permitted decisions completed, failures recovered and providers replaced. Membership counts, standards counts and graph size are useful activity measures only when their limits are acknowledged. They are not substitutes for realized industrial value.

Europe's strategic contribution should be an internationally usable capacity for cooperation. Jointly maintained profiles, portable evidence and replaceable services can let partners pursue industrial and public-interest objectives without requiring one platform or one jurisdiction to control every relationship. Participating regions should help govern the interfaces and contribute implementations. The counter-model becomes credible when participants retain useful choices across changes of vendor, operator and circumstance. Durable agency is the objective; demonstrated continuity and fair participation are the tests.

13 Scenarios limitations and the research agenda

Four scenarios clarify the uncertainty. In **platform-led accommodation**, integrated vendors improve exports, interfaces and profile support; customers gain useful interoperability while substantial operational value remains within platforms. In **cooperative federation**, implementations in several regions preserve selected guarantees, participants share governance and viable firms maintain the infrastructure. In **nominal federation**, open components depend on a few indispensable gateways, registries or consulting integrations. In **regional fragmentation**, technically similar systems diverge through incompatible profiles, restrictions or exclusive trust arrangements, increasing the cost of cross-border cooperation. These are analytical scenarios, not probability forecasts.

Discriminating observations include successful cross-vendor substitution, the share of decision-relevant claims covered by tested profiles, time and cost to onboard, continuity during the loss of an optional intermediary, and net value after governance and verification costs. The international dimension adds participation by regional suppliers, transparent decisions on issuer acceptance, conformance in both directions and the cost of justified jurisdictional differences. A successful technical exchange should be reported separately from institutional acceptance or legal recognition. A further scenario stress test should remove a selected cloud or model service and assess what remains operational. That infrastructure-level test is proposed future work; it is not included in the current fifteen-case valve fixture.

The documentary comparison has several limitations. Provider documentation is evidence of specified mechanisms, not independently reproduced performance. Industry associations report their own participation and benefit measures. The corpus is targeted rather than comprehensive.

The synthetic case is intentionally narrow and cannot establish functional safety, universal legal acceptance or the economics of large-scale deployment. The initial implementation candidates share important dependencies.

The theoretical extension is also conditional. Reusable identity and semantics may increase opportunities for cooperation, but translation losses, weak incentives, untrusted measurements and expensive governance can prevent value realization. A captured issuer can produce a cryptographically valid false claim. A technically replaceable service may remain economically dominant through scale or expertise. These conditions belong in the research agenda rather than being excluded from the counter-model's evaluation.

The next studies should test complete contracts, then independent implementations and realistic operating conditions. Legal analysis should examine specific B2B and B2G use cases. Economic work should compare actual lifecycle costs and realized benefits, with appropriate counterfactuals. Safety-critical and physical-AI deployments need their own engineering arguments. This sequence converts a strategic proposition into evidence without requiring the first prototype to establish every claim at once.

14 Conclusion

Palantir demonstrates the value of making an ontology operational: connecting data and domain meaning to logic, governed actions and applications. A global cooperation model should preserve that depth while allowing participants to combine capabilities across independently governed implementations. The proposed meta-platform makes agreed meaning, verifiable evidence and authorized requests reusable beyond the environment that first assembled them.

Europe's dataspace work is a first step that now requires **fast-track modernisation** for AI-first operation. It must not be treated as a proven infrastructure foundation for IPCEI-AI. Its strategic opportunity is to convene partners around substantial engineering work: meet declared throughput, latency and scale requirements while preserving access control, verifiability and participant authority. Documented Japanese interoperability work, Chinese standards-cooperation initiatives, US contributions to open web standards and multilateral traceability efforts provide reasons for global coordination. They differ in scope and maturity; together they establish opportunities for collaboration rather than a completed global system.

Access-controlled, verifiable evidence graphs built on open international standards and industry-agreed semantics are the proposed foundation. They can support physical products, virtual AI services and physical AI when combined with domain-specific assurance and local authority. International trust emerges from attributable claims, accepted responsibilities, accountable governance and tested behavior. Cryptographic validity alone does not establish truth, safety or legal recognition.

The constructive counter-model to closed platform dependence is therefore sovereignty through cooperation. Europe can connect its industrial coordination and digital diplomacy with the innovation and requirements of global partners, including US innovators and appropriate engagement across Asia-Pacific and China. Shared interfaces can support competition and resilience while participants retain control over permitted data use and consequential actions. The next test is to demonstrate these properties in jointly governed cases, preserve them through provider changes, and report the costs and failures as carefully as the successes.

15 Data availability and disclosures

The accompanying local versioned package contains the manuscript, editable figures, source records, frozen architecture baselines, component candidates, synthetic fixture and unexecuted protocol. A public archival release or DOI has not been created. The author is affiliated with Spherity GmbH and is a co-author of the meta-platform paper. He is listed as a facilitator of the W3C DPP and Business Wallets Task Force. Spherity is represented among the credited contributors to the Recognized Entities draft. These roles are disclosed as relevant interests. The original figures are analytical syntheses. This is an editorially reviewed research draft, not external peer review, certification or endorsement by the organizations discussed. [99] [101]

A Common blocks and detailed local flows

The functional IDs remain stable across P1 and P2. B10 is intentionally separate from B04 so the vocabulary is not mistaken for the populated graph. The ontology runtime combines B05, B06 and B08; B10 overlaps its write/indexing path. G01 and G02 apply to all blocks and exchanges. The core figures show principal flows; the following table specifies the shared responsibilities without prescribing one product per block.

Table 1: Common functional blocks and implementation responsibilities.

Block	**Responsibility**
B01 Source access	Obtain permitted source content and remote services.
B02 Preparation	Extract and normalize; retain context and uncertainty.
B03 Data substrate	Retain assets, snapshots and permitted remote references.
B04 Ontology model and vocabulary	Define types, relationships, capabilities and approved meanings.
B10 Knowledge Graph Builder	Map, attribute, validate, index and refresh instances and links.
B05 Operational graph and read services	Serve authorized objects, graph views and query results.
B06 Logic and model execution	Execute declared rules, calculations, models or entailment.
B07 AI orchestration	Coordinate permitted tools and bounded delegated tasks.
B08 Action execution	Enforce current preconditions; edit, invoke, receipt and reconcile.
B09 Applications	Support human interaction, explanations and approvals.
G01 Governance	Identity, authority, policy, evidence admission and audit.
G02 Lifecycle	Versioning, evaluation, promotion, observation and recovery.

Table 2: Corresponding local flow groups in both architecture patterns.

Flow	**Endpoints and purpose**
F01 / L01	B01 → B02. Source content / changes.
F02 / L02	B02 → B03. Prepared datasets / media.
F03 / L03	B03 → B10. Backing data / updates.

Flow	Endpoints and purpose
F04 / L04	B04 ↔ B10. Definitions to builder; configured metadata edits back.
F05 / L05	B04 → B05, B06, B08. Types, functions and action configuration.
F06 / L06	B10 → B05. Index / refresh objects and links.
F07 / L07	B08 → B10. Action edits enter the indexing path.
F08 / L08	B09 ↔ B07. User request / clarification / answer.
F09 / L09	B09 ↔ B05. Query / results.
F10 / L10	B07 ↔ B05. Query / retrieved context.
F11 / L11	B09 ↔ B06. Function call / result.
F12 / L12	B07 ↔ B06. Tool call / result.
F13 / L13	B08 ↔ B06. Configured function call / edits or result.
F14 / L14	B06 ↔ B05. Object query / operational context.
F15 / L15	B09 ↔ B08. Action request / execution result.
F16 / L16	B07 ↔ B08. Permitted tool call / execution result.
F17 / L17	B08 ↔ X02. External call / response (separate side effect).
F18 / L18	X01 ↔ B01. Selected data access / exchange.
F19 / L19	X02 → B01. Subsequent observed state / events.
F20 / L20	X03 ↔ B05, B06, B08. Configured query / action calls and results.
F21 / L21	X04 ↔ B06, B07. Configured model calls / results.
F22 / L22	B01 ↔ B03. Virtual / direct access route.
F23 / L23	B08 ↔ B05. Read transaction context / result.
F24 / L24	B07 ↔ B04, B10. Governed engineering proposals; review before promotion.
F25 / L25	TRIGGER → B06, B07, B08. Schedule, event or API dispatch; LLM optional.

The paired flow IDs F01–F25 and L01–L25 share functional purposes. A bidirectional group includes different request and return payloads; it does not grant equal authority in both directions. Endpoint arrays represent configured functions or alternatives rather than a Cartesian product of connections. Implementation-specific error messages, ordering and retry protocols remain outside this abstract register.

B Federation exchange and dependency contracts

P2 adds three metadata-supply groups and nine governed exchange groups. The versioned flow register and exchange-contract register contain the complete X01–X12 mapping and contract fields. Table 3 summarizes the obligations an implementation must make concrete. A selected protocol may distribute these responsibilities across several services.

Table 3: Boundary obligations for the proposed federation.

Exchange responsibility	Required declaration or check
Shared semantic release	Publisher authority, version, dependency manifest, integrity and accepted profile.
Catalog and resolver access	Permitted metadata, endpoint identity, scope and freshness.
Trust and status supply	Accepted issuer scope, authority, key bindings, status age and outage policy.
Identity and delegation	Organization, presenting actor, audience, task, operation and delegated authority.
Policy and agreement	Supported mandatory terms, purpose, disclosure, obligations and conflict behavior.
Evidence disclosure	Exact protected payload, subject, original issuer and admitted graph projection.
Remote query or computation	Snapshot, provenance, completeness and verification level of the result.
Action and outcome	Separate local authorization, current preconditions, idempotency and receipt.
Change and reassessment	Profile/status dependencies, affected decisions and retained historical context.

Approved local caches require authenticated source issue times and explicit validity windows. Cache residence time is not authority freshness. Dependency manifests should cover imported contexts, vocabularies, shapes and trust material. These are design requirements; TUF-inspired semantic publication, policy-composition behavior and universal cross-connector enforcement are not reported as completed implementations.

C Candidate components and planned experiment

Candidate A proposes Jena 6.2.0, while Candidate B proposes RDF4J 6.1.0 with separate application adapters. The component map records their different Java requirements and a release-catalog discrepancy for RDF4J. EDC Connector and IdentityHub 0.18.1 are observed stable candidates, while 1.0.0-RC5 is separately identified as pre-release. These selections require compatible distribution and protocol tests before deployment; they are not a tested dependency lock. [79] [80] [81]

The target exchange versions are DSP 2025-1-err2 and DCP v1.0.1. The VC library candidate, its securing-format limits and the selected UNTP profile must be reconciled before proofs are generated. Cryptographic primitives alone are not a complete verifier of authority, semantics, status and intended use. The initial experiment does not test independent implementations of all these components. [84] [85] [82] [83]

The fixed decision time is 1 November 2026 at 10:00 UTC. Observation age is measured from the end of the laboratory interval. Maximum evidence age is 72 hours, maximum authenticated status age is 15 minutes and maximum authority-bundle age is 24 hours. These are synthetic experiment limits. Protected payloads, issuer keys, profile versions, policy mappings and clocks must be pinned before execution. The fixture contains preparation instructions and oracle values, not generated cryptographic proofs.

Table 4: Planned test cases and expected disposition.

Case	Expected behavior
C01 Healthy fixture	Eligible; one separately approved inventory transition.
C02 Missing evidence	Hold without claiming the test does not exist anywhere.
C03 Denied access	Hold; no protected disclosure.
C04 Contradictory claims	Hold; preserve conflicting attribution.
C05 Revoked evidence	Hold under the accepted status policy.
C06 Wrong authority scope	Hold despite a potentially valid signature.
C07 Unsupported mandatory policy	Hold before treating the term as enforced.
C08 Stale status during outage	Hold; current status is unestablished.
C09 Unsupported remote conclusion	Hold; embedded instructions remain untrusted data.
C10 ERP commit with lost response	Reconcile; exactly one new commit.
C11 Consumed request replay	No new decision or commit; retain prior state.
C12 Replacement provider	Eligible if the same contract and new audience binding hold.
C13 Tampered protected payload	Hold on integrity failure.
C14 Unapproved semantic profile	Hold pending approved interpretation.
C15 Valid measurement above limit	Engineering rejection; no release commit.

Each case has three reset functional repetitions per pattern. C11 explicitly begins from a prior successful state. C10 retries with a fresh nonce and the same operation digest and idempotency key; it does not reuse a consumed signed request. C12 changes the serving implementation and authorized operator without automatically changing the accepted laboratory authority. Record replacement effort separately from initial integration.

Figure 9 aligns the same eight-stage journey for both patterns. The machine-readable protocol defines permitted outputs, confidential canaries, side-effect counts and held-out ground truth. Only after functional correctness are thirty cold and thirty warm healthy performance runs proposed, following excluded preflight runs. Record resources, network conditions and caches that cannot be reset; unmatched deployment resources do not support a causal performance ranking.

All cases remain planned and not run. Future extensions include timed revocation-propagation experiments, independent connector and verifier substitution, infrastructure-provider outage tests and domain-specific safety or legal assessment. Those extensions must receive their own fixtures, threat model and denominators before results are claimed.

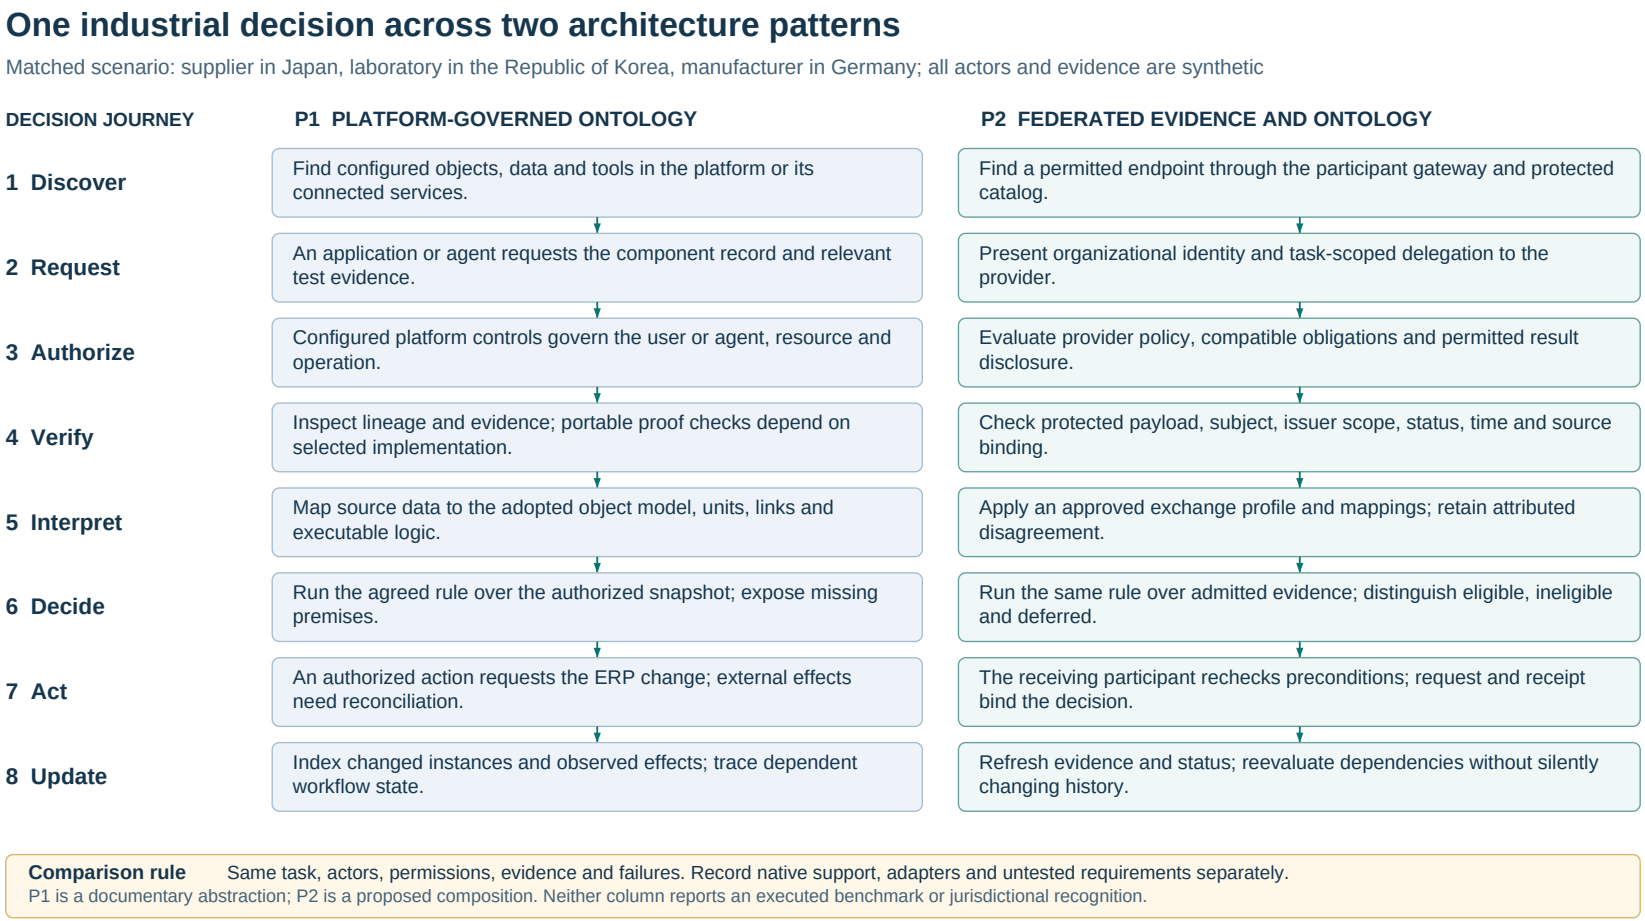


Figure 9: The same discovery, request, authorization, verification, interpretation, decision, action and update responsibilities are assessed in both patterns. The fixture and its oracle are synthetic. This diagram and the test-case table define a proposed experiment; they do not report comparative execution results.

References

[1] Palantir. Ontology system. Undated / rolling source. Accessed 1 November 2026.

[2] European Commission, DG CONNECT. Common European data spaces. Undated / rolling source. Accessed 1 November 2026. Updated 2026-08-14.

[3] Michael Franklin, Alon Halevy, David Maier. From Databases to Dataspaces: A New Abstraction for Information Management. 34(4), pp. 27–33. Published 2005. Accessed 1 November 2026. Inspected full text.

[4] Dr. Carsten Stöcker. Evidence Graphs for Industrial AI. Published 2026-07-09. Accessed 1 November 2026. Updated 2026-09-07. Authoritative full paper.

[5] European Commission, High Representative of the Union for Foreign Affairs and Security Policy. An International Digital Strategy for the European Union. JOIN(2025) 140 final. Council transmission 10593/25 dated 19 June 2025. Published 2025-06-05. Accessed 1 November 2026.

[6] European Commission, Directorate-General for Digital Services. Cloud Sovereignty Framework: Implementation guidance. Published 2026-06-01. Accessed 1 November 2026. Framework explanation. Procurement record.

[7] Michael Shea, Samuel M. Smith, Carsten Stöcker. Decentralized Identity as a Meta-platform: How Cooperation Beats Aggregation. Published 2019-11-19. Accessed 1 November 2026. Version 1.6. Inspected Markdown v1.6; PDF byte identity not verified.

[8] Michael L. Katz, Carl Shapiro. Systems Competition and Network Effects. Journal of Economic Perspectives, 8, 2, 93–115. DOI: 10.1257/jep.8.2.93. Published 1994. Accessed 1 November 2026. Author hosted full text.

[9] Samuel M. Smith. Meta-Platforms and Cooperative Network-of-Networks Effects: Why Decentralized Platforms Will Eat Centralized Platforms. Published 2019-03-25. Accessed 1 November 2026.

[10] Leo Van Hove. Testing Metcalfe's law: pitfalls and possibilities. Information Economics and Policy, 37, 67–76. DOI: 10.1016/j.infoecopol.2016.09.001. Published 2016-12. Accessed 1 November 2026.

[11] Aidan Hogan et al. Knowledge Graphs. 54(4), Article 71, pp. 1–37. DOI: 10.1145/3447772. Author version 6, 11 September 2021. Published 2021. Accessed 1 November 2026.

[12] Shirui Pan et al. Unifying Large Language Models and Knowledge Graphs: A Roadmap. DOI: 10.1109/TKDE.2024.3352100. Author preprint version 3, 25 January 2024. Published 2024. Accessed 1 November 2026.

[13] W3C. OWL 2 Primer, second edition. Published 2012-12-11. Accessed 1 November 2026.

[14] W3C. Shapes Constraint Language (SHACL). Published 2017-07-20. Accessed 1 November 2026.

[15] W3C. RDF 1.1 Concepts and Abstract Syntax. Published 2014-02-25. Accessed 1 November 2026.

[16] W3C. JSON-LD 1.1. Published 2020-07-16. Accessed 1 November 2026.

[17] GS1. Core Business Vocabulary, release 2.0. Published 2022-06. Accessed 1 November 2026.

[18] Palantir. Multi-organisation ecosystems. Undated / rolling source. Accessed 1 November 2026.

[19] Palantir. Ontology backend. Undated / rolling source. Accessed 1 November 2026.

[20] Palantir. Pipeline input datasets. Undated / rolling source. Accessed 1 November 2026.

[21] Palantir. Excel connector. Undated / rolling source. Accessed 1 November 2026.

[22] Palantir. Document processing. Undated / rolling source. Accessed 1 November 2026.

[23] Palantir. Document Intelligence concepts. Undated / rolling source. Accessed 1 November 2026.

[24] Palantir. Multimodal Data Plane. Undated / rolling source. Accessed 1 November 2026.

[25] Palantir. Ontology building. Undated / rolling source. Accessed 1 November 2026.

[26] Palantir. Types reference. Undated / rolling source. Accessed 1 November 2026.

[27] Palantir. Object type creation. Undated / rolling source. Accessed 1 November 2026.

[28] Palantir. Link type creation. Undated / rolling source. Accessed 1 November 2026.

[29] Palantir. AIP architecture. Undated / rolling source. Accessed 1 November 2026.

[30] Palantir. Chatbot tools. Undated / rolling source. Accessed 1 November 2026.

[31] Palantir. Ontology-augmented generation. Undated / rolling source. Accessed 1 November 2026.

[32] Palantir. Action consistency and isolation. Undated / rolling source. Accessed 1 November 2026.

[33] Palantir. AI FDE overview. Undated / rolling source. Accessed 1 November 2026.

[34] Palantir. AIP Evals. Undated / rolling source. Accessed 1 November 2026.

[35] Palantir. B2B collaboration. Undated / rolling source. Accessed 1 November 2026.

[36] Palantir. Ontology export and import. Undated / rolling source. Accessed 1 November 2026.

[37] Palantir. Current interoperability architecture. Undated / rolling source. Accessed 1 November 2026.

[38] Palantir. Foundry connector. Undated / rolling source. Accessed 1 November 2026.

[39] Eclipse EDC. EDC Adopters Manual: Connector, Control Plane, Data Plane and Identity Hub. Undated / rolling source. Accessed 1 November 2026.

[40] International Data Spaces Association. IDSA Rulebook: What Is a Dataspace?. Undated / rolling source. Accessed 1 November 2026.

[41] Sebastian Steinbuss et al. Semantic Interoperability in Data Spaces. . DOI: 10.5281/zenodo.10964377. Published 2024-04. Accessed 1 November 2026. Version 1.0.

[42] W3C. Data Catalog Vocabulary (DCAT), version 3. Published 2024-08-22. Accessed 1 November 2026.

[43] W3C. Verifiable Credentials Data Model v2.0. Published 2025-05-15. Accessed 1 November 2026.

[44] Object Management Group. Structured Assurance Case Metamodel (SACM), version 2.3. Published 2023-10. Accessed 1 November 2026.

[45] W3C. PROV-O: The PROV Ontology. Published 2013-04-30. Accessed 1 November 2026.

[46] W3C. Verifiable Credential Data Integrity 1.0. Published 2025-05-15. Accessed 1 November 2026.

[47] W3C. SPARQL 1.1 Federated Query. Published 2013-03-21. Accessed 1 November 2026.

[48] W3C. Bitstring Status List v1.0. Published 2025-05-15. Accessed 1 November 2026.

[49] W3C. ODRL Information Model 2.2. Published 2018-02-15. Accessed 1 November 2026.

[50] Eclipse EDC. EDC Policy Engine. Undated / rolling source. Accessed 1 November 2026.

[51] Scott Rose, Oliver Borchert, Stu Mitchell, Sean Connelly. NIST SP 800-207: Zero Trust Architecture. Published 2020-08. Accessed 1 November 2026.

[52] The Update Framework. TUF specification and python-tuf releases. Undated / rolling source. Accessed 1 November 2026. Python TUF implementation releases.

[53] IETF. RFC 9111: HTTP Caching. Published 2022-06. Accessed 1 November 2026.

[54] Dylan Yaga, Peter Mell, Nik Roby, Karen Scarfone. NISTIR 8202: Blockchain Technology Overview. Published 2018-10. Accessed 1 November 2026.

[55] UK Competition and Markets Authority. Cloud infrastructure services: Final decision report. Decision dated 31 July 2025; full report listed 1 August 2025. Dated 2025-07-31. Accessed 1 November 2026. Investigation and final decision record.

[56] UK Competition and Markets Authority. Update on cloud actions: covering March to October 2026. Published 2026-10-16. Accessed 1 November 2026. Business software and cloud services. CMA announcement.

[57] European Commission, DG CONNECT. Data Act explained. Undated / rolling source. Accessed 1 November 2026. Updated 2025-12-15.

[58] European Commission, DG COMP and DG CONNECT. Commission reaches preliminary position that Amazon's and Microsoft's market leading cloud services should be designated under the DMA. Published 2026-06-25. Accessed 1 November 2026. Related primary source 1. Related primary source 2.

[59] European Data Protection Board. Guidelines 02/2024 on Article 48 GDPR, version 2.1. Adopted 2025-06-04. Dated 2025-06-20. Accessed 1 November 2026. Guidelines and version record.

[60] Spherity GmbH / Dr. Carsten Stöcker. PACE: Verifiable Safety and Cyber Evidence Graphs for Physical AI. Published 2026-09-09. Accessed 1 November 2026.

[61] Spherity GmbH. MiPA Assurance Viewer. Undated / rolling source. Accessed 1 November 2026.

[62] Spherity GmbH / Dr. Carsten Stöcker. Governed Memory for Physical AI: Evidence Roles, Integrity and Revocation from Experience to Action. Published 2026-10-20. Accessed 1 November 2026. Updated 2026-10-21.

[63] European Parliament and Council / EUR-Lex. Regulation (EU) 2024/1781 — Ecodesign for Sustainable Products Regulation. Dated 2024-06-13. Accessed 1 November 2026.

[64] ISO; UK Health and Safety Executive. Industrial robot safety and functional safety scope. Dated 2025-02. Accessed 1 November 2026. HSE functional safety guidance.

[65] European Parliament and Council / EUR-Lex. Regulation (EU) 2023/1230 on machinery. Dated 2023-06-14. Accessed 1 November 2026.

[66] European Commission AI Act Service Desk; EUR-Lex. AI Act current classification and scope. Dated 2026-07-27. Accessed 1 November 2026. AI Act current Article 2. AI Act legislative text.

[67] European Parliament and Council; European Commission. Regulation (EU) 2026/1744 — Digital Omnibus on AI. Published 2026-07-24. Accessed 1 November 2026. Commission entry-into-force notice. Commission implementation timetable. Council legislative text.

[68] European Parliament and Council / EUR-Lex; Medical Device Coordination Group. Medical Devices Regulation and intended-purpose guidance. Dated 2017-04-05. Accessed 1 November 2026. MDCG medical software guidance.

[69] European Parliament and Council / EUR-Lex. Regulation (EU) 2024/2847 — Cyber Resilience Act. Dated 2024-10-23. Accessed 1 November 2026.

[70] European Commission DG Trade; European Parliament and Council / EUR-Lex. EU dual-use export controls. Dated 2021-05-20. Accessed 1 November 2026. EU dual-use regulation.

[71] European Commission, Simpl Programme. Simpl-Open development update: Gdynia available and Hamburg preparations underway. Undated / rolling source. Accessed 1 November 2026.

[72] Factory-X project consortium. Factory-X: The digital ecosystem. Undated / rolling source. Accessed 1 November 2026.

[73] Catena-X Automotive Network e.V. Catena-X network grows by more than a factor of five to over 1000 companies. Published 2026-09-21. Accessed 1 November 2026.

[74] European Commission, DG CONNECT. Commission welcomes the design of first Important Project of Common European Interest in AI. Published 2026-09-16. Accessed 1 November 2026. Updated 2026-09-17.

[75] Ireland, Department of Enterprise, Tourism and Employment. Ireland backs major European AI initiative to strengthen innovation, investment and technological resilience. Published 2026-10-30. Accessed 1 November 2026.

[76] Information-technology Promotion Agency, Japan, and Catena-X Automotive Network e.V. Catena-X and Ouranos Ecosystem Successfully Demonstrate Data Space Interoperability. Published 2025-03-31. Accessed 1 November 2026.

[77] National Data Administration, China. Digital China Development Report (2025) [in Chinese]. Report dated 2026-05. Accessed 1 November 2026.

[78] UNECE / UN/CEFACT UNTP project. UN Transparency Protocol: FAQ. Undated / rolling source. Accessed 1 November 2026.

[79] Apache Software Foundation. Apache Jena release and graph/validation documentation. Undated / rolling source. Accessed 1 November 2026. Project overview. SHACL documentation.

[80] Eclipse RDF4J. Eclipse RDF4J release, SHACL and FedX documentation. Undated / rolling source. Accessed 1 November 2026. SHACL validation documentation. Federated query documentation. Eclipse project catalog.

[81] Eclipse EDC. Eclipse EDC releases and dataspace profiles. Undated / rolling source. Accessed 1 November 2026. IdentityHub releases. Dataspace profiles documentation.

[82] Digital Bazaar. Digital Bazaar Verifiable Credentials JS library and changelog. Undated / rolling source. Accessed 1 November 2026. Library changelog.

[83] UNECE / UNTP. UNTP current technical status and versioned VC profile. Undated / rolling source. Accessed 1 November 2026. Version 0.7.0 credential profile. Current credential profile.

[84] Eclipse. Dataspace Protocol 2025-1, errata 2. Undated / rolling source. Accessed 1 November 2026.

[85] Eclipse. DCP v1.0.1 release and trust model. Undated / rolling source. Accessed 1 November 2026. Protocol trust model.

[86] Council of the European Union. Council approves conclusions on a strategy for the EU's engagement in global digital affairs. Published 2025-11-20. Accessed 1 November 2026. Updated 2025-11-24.

[87] European Commission. European Data Union Strategy. Undated / rolling source. Accessed 1

November 2026. Updated 2026-05-18.

[88] United Nations Centre for Trade Facilitation and Electronic Business (UN/CEFACT). UNTP Core Specification. Undated / rolling source. Accessed 1 November 2026.

[89] United Nations Economic Commission for Europe (UNECE), UN/CEFACT. Recommendation No. 46: Enhancing traceability and transparency of sustainable value chains in the garment and footwear sector. ECE/TRADE/463. Published 2022-01. Accessed 1 November 2026. Report of the United Nations Centre for Trade Facilitation and Electronic Business on its twenty-seventh session. Enabling Sustainability and Circularity in the Garment and Footwear Sector: Policy Developments and Industry Perspectives on Traceability and Transparency.

[90] United Nations Centre for Trade Facilitation and Electronic Business (UN/CEFACT). Report of the United Nations Centre for Trade Facilitation and Electronic Business on its Thirty-First Session. ECE/TRADE/C/CEFACT/2025/2. Published 2025-08-06. Accessed 1 November 2026. Recommendation No. 49 (ECE/TRADE/493): Transparency at Scale: Fostering Sustainable Value Chains.

[91] UN Transparency Protocol project / UN/CEFACT. Governance. Undated / rolling source. Accessed 1 November 2026. Documentation Versions. Sovereignty and Jurisdiction.

[92] Standards Australia. New Universal Data Protocol for the built environment announced at COP29. Published 2024-11-16. Accessed 1 November 2026. Climate Reporting.

[93] Information-technology Promotion Agency, Japan; joint release with Data Society Alliance, Robot Revolution & Industrial IoT Initiative, and The University of Tokyo. Agreement to Jointly Promote the Dataspace Technical Concept As Open Data Spaces. Published 2025-10-15. Accessed 1 November 2026.

[94] National Data Administration, China. Trusted Data Spaces Development Action Plan (2024-2028) [in Chinese]. Issued 21 November 2024. Author-translated title. Published 2024-11-23. Accessed 1 November 2026. Official action-plan attachment.

[95] National Data Administration, China; source credited to the National Technical Committee on Data Standardization. SAC/TC609 and the International Data Spaces Association sign a memorandum of understanding in Beijing [in Chinese]. Reports signature on 17 December 2025. Author-translated title. Published 2025-12-23. Accessed 1 November 2026.

[96] National Development and Reform Commission, China, Department of High-Tech Industry. Action Plan for Cooperation and Development in Artificial Intelligence [in Chinese]. Author-translated title. Published 2026-07-17. Accessed 1 November 2026.

[97] U.S. Small Business Administration, SBIR/STTR Award Data. Verifiable Credentials for First Responders. DHS SBIR Phase II award, contract 70RSAT24C00000010. Scheduled 8 April 2024 to 7 April 2026; results not inferred from schedule. Undated / rolling source. Accessed 1 November 2026.

[98] World Wide Web Consortium (W3C). Verifiable Credentials Working Group Charter. Charter period 11 March 2026 to 31 March 2028; tentative deliverables. Undated / rolling source. Accessed 1 November 2026.

[99] W3C. DPP and Business Wallets Task Force. Undated / rolling source. Accessed 1 November 2026.

[100] W3C task-force contributors. DPP and Business Wallet use cases and requirements. Undated / rolling source. Accessed 1 November 2026.

[101] W3C. Recognized Entities v1.0. Published 2026-09-06. Accessed 1 November 2026.