A European Answer to Palantir

A globally collaborative counter model for industrial AI

Executive brief 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.

Executive summary

Europe's dataspaces are an important first step. Their conceptual roots predate the current frontier-AI and large-language-model wave. Experience in controlled sharing is valuable, but it does not establish that existing implementations meet AI-first requirements for high throughput, low latency and large data volumes. Europe should pursue **fast-track modernisation**, with international partners, rather than assume that existing exchange arrangements are ready for IPCEI-AI ambitions.

Palantir shows the value of an operational ontology connecting data, meaning, permissions and actions. Europe's opportunity is comparable operational depth across independently governed systems, with usable provider choice. Companies should process their own records, exchange permitted contributions and connect the evidence needed for a particular task.

The foundation remains **access-controlled, verifiable evidence graphs built on open international standards and industry-agreed semantics**: connected claims with shared meaning, checkable origins and clear access rules. Modernisation must meet the required performance and trust guarantees together, preserving access control, verifiability and participant authority. This requires engineering and investment; it is not a promise of cost-free verification. Bulk data and evidence about that data can follow different, coordinated paths.

A proposed 90-day programme should test a bounded industrial workflow under realistic load, including denial, failure and provider replacement. Measure throughput, response times and total cost alongside evidence quality and unauthorized disclosure. Policymakers should fund shared vocabulary maintenance, compatible interfaces and independent tests with partners in the United States, Asia-Pacific and beyond. The strategic offer is globally usable cooperation infrastructure with accountable local control. Its readiness must be demonstrated.

Keywords. industrial AI; digital sovereignty; Palantir; dataspaces; ontology engines; knowledge graphs; digital trust diplomacy; international cooperation

Contents

1 The geopolitical choice is who controls industrial cooperation **3**

2 Why the ontology is central to Palantir **3**

3 What knowledge graphs can improve for AI **5**

4 A humanoid robot makes the evidence relationships visible **7**

5 An ecosystem of ontology engines **11**

6 Fast-track modernisation for industrial AI **11**

7 Digital trust diplomacy in practice **12**

8 A different business model needs collaborative leadership **15**

9 A practical first mover programme **16**

10 An invitation to build the next operational layer together **17**

11 Scope and disclosure **17**

References **18**

1 The geopolitical choice is who controls industrial cooperation

As AI moves from answering questions to coordinating operations, the systems underneath it acquire strategic importance. They determine which suppliers are visible, what a compliance claim means, whose test result is accepted and which agent may request a physical action. A company that cannot carry these capabilities to another provider may retain its files while losing practical control over its operations.

Europe needs an operational alternative to platform dependence. A policy focused only on where data is hosted will miss part of the problem. Definitions, mappings, access rules, software integrations and evidence histories can be harder to replace than storage. In a geopolitical disruption, commercial dispute or service outage, these dependencies can limit the choices available to industry and public authorities.

The concern is concrete. In March 2026, the UK Competition and Markets Authority described cloud switching and interoperability barriers, alongside software licensing concerns, following its cloud market investigation. Its October update reported progress on fees and interoperability, while identifying further work on switching and identity and access management. These are findings about cloud markets, not findings against Palantir. They illustrate why technical and commercial exit options deserve attention before dependence becomes entrenched. [1] [2]

Palantir is a useful and demanding benchmark because it turns an integrated platform into an operating environment for decisions. Our counter-model addresses the platform-centred business model directly: the operator of a useful service should compete on its value, while the shared language of an industry and the evidence needed to participate remain usable across providers. Europe can pursue that objective with US innovators and partners across Asia-Pacific. Nationality alone does not establish openness, and a European platform can create the same dependencies.

This is a strategic proposal, not a description of an already completed European system. Europe's advantage would come from organizing cooperation at industrial scale: interoperable services, credible supplier choice and shared institutions that participants can influence. Fragmented national pilots, incompatible vocabularies or mandatory dependence on a single European hub would weaken that offer.

2 Why the ontology is central to Palantir

An **ontology** defines the things an organization works with and how they relate: a component, a test, a supplier, a work order and a permitted action. A **knowledge graph** populates those definitions with linked records. An **ontology engine** makes the model usable through queries, calculations, permissions and operations. The vocabulary supplies the agreed terms; the engine connects those terms to work.

Palantir's contribution is the integration of these capabilities. Its documentation puts data, logic, actions and security together, with an Ontology Language, Engine and Toolchain. An AI assistant can work with a defined component and approved tools, rather than repeatedly reconstructing the business from loose text. This is a core architectural capability. It does not, by itself, explain all of Palantir's commercial success. [3]

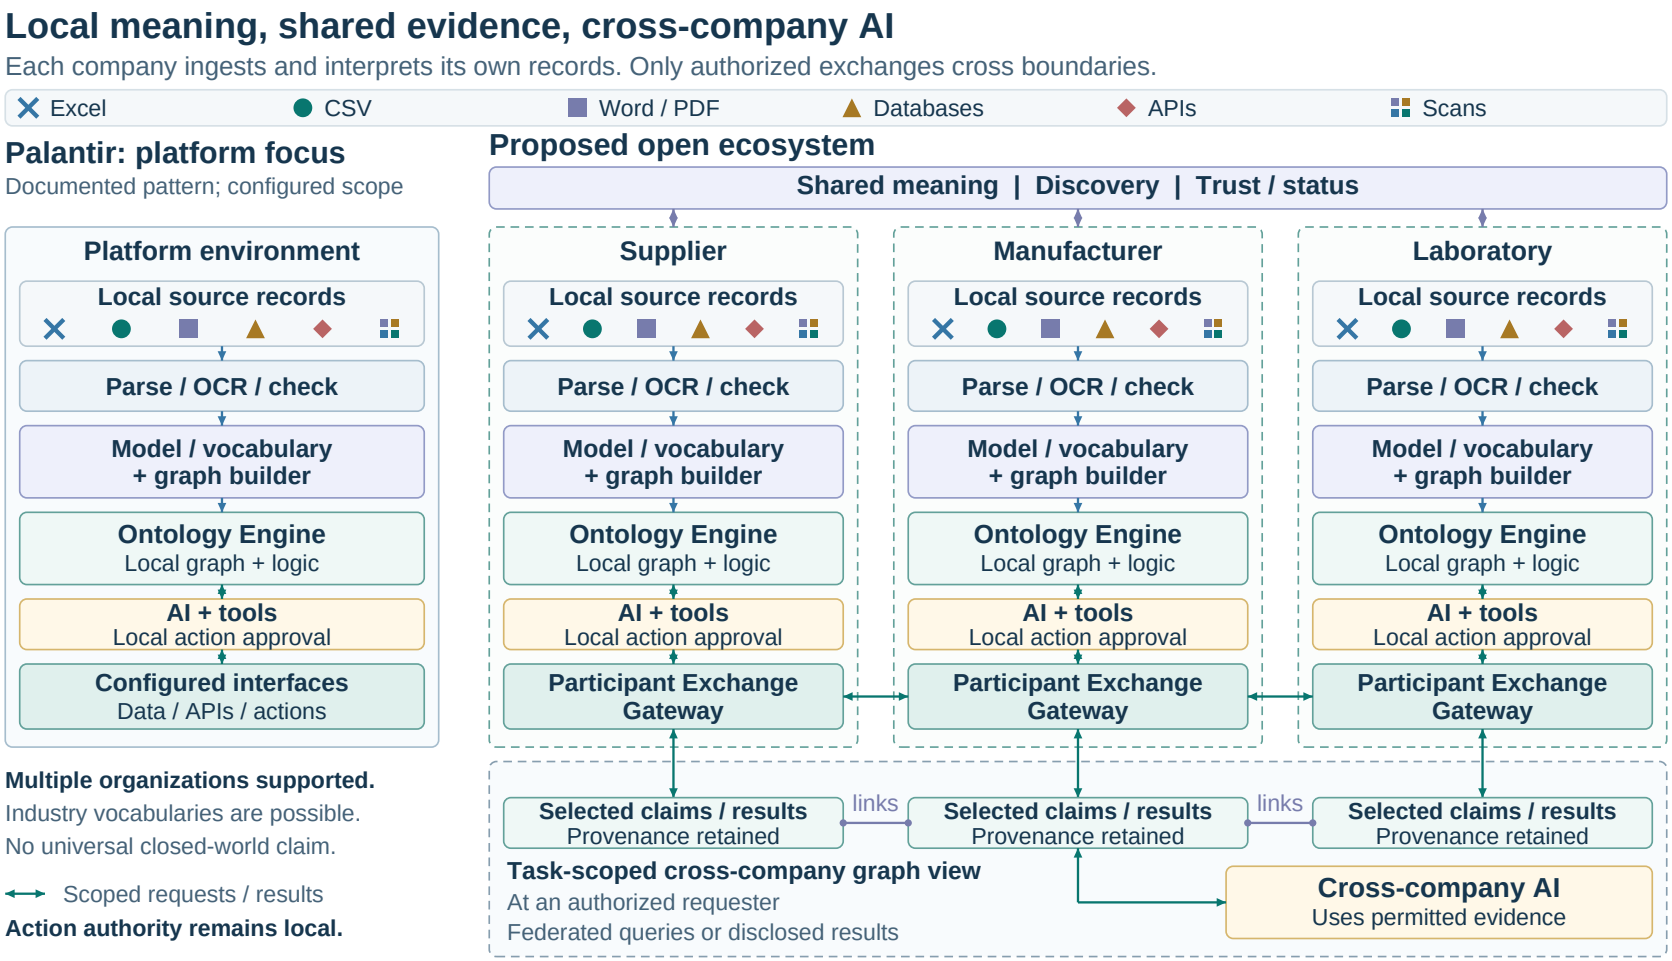


Figure 1: Each participant ingests, checks and links its own records. Its gateway exchanges authorized requests, evidence and results. An authorized requester composes selected contributions into a cross-company graph view, preserving sources and use conditions. Shared meaning, discovery and trust serve all participants. This is a task-specific view, not a pooled database. Local authority controls actions. The legend identifies formats, not trust. Palantir supports multiple organizations; the comparison concerns governance across implementations. Proposed architecture, not measured performance.

The path begins with ordinary industrial information: Excel sheets, CSV files, Word documents, PDFs, databases, application interfaces and scanned records. These inputs must be accessed, parsed and checked. Optical character recognition, or OCR, turns a scan into candidate text; it does not certify that the extracted value is correct. Mapping then connects the information to identifiers, units and business concepts. Palantir documents support for diverse data modalities and separate document-processing workflows. [4] [5]

Figure 1 follows three steps: **process locally, exchange under control, then connect the permitted evidence**. Each company's knowledge graph builder links records using its business model and mappings to shared industry terms. Exchange can carry selected claims, evidence or query results; it need not transfer the underlying files or the whole local graph. The resulting view connects only the information available and permitted for that task, keeping missing evidence visible. Here, model means the ontology's business model unless an AI model is explicitly named. Building a knowledge graph is different from training a language model.

Palantir already supports multi-organization ecosystems, participant-specific controls and common ontologies that can use industry standards. The fair comparison is platform-centred integration versus cooperation across independently governed implementations. We do not claim that Palantir only serves one company, cannot interoperate, or universally applies formal closed-world logic. [6]

3 What knowledge graphs can improve for AI

Consider a manufacturer replacing a valve. The supplier's spreadsheet contains a serial number, the laboratory's PDF records a leakage measurement, and a database holds the manufacturer's acceptance limit. A scan may contain an older report for a different valve. Finding similar words across these files is insufficient. The decision depends on the right component, the right test, compatible units and the applicable limit.

A well-maintained graph makes those relationships explicit. It can connect the exact valve to the relevant test and preserve the report, its issuer, its date and the transformation that extracted the measurement. A query can retrieve that connected evidence. A deterministic calculation can compare the measurement with the limit. The AI can explain the result and cite the supporting records. Figure 2 shows how this division of work reduces the amount of interpretation left to a language model.

Fewer tokens can mean lower operating cost. Tokens are the fragments of text processed by an AI model. Selecting the relevant facts can avoid repeatedly sending hundreds of pages. That can reduce prompt length and irrelevant context. The saving must include the cost of constructing, maintaining and querying the graph. Microsoft's GraphRAG research reports large query-context savings against summarizing an entire source corpus; that comparison is not a universal saving against ordinary document retrieval. [7]

Better grounding can reduce unsupported answers. The system can require a source for each material claim and reject a suggested relationship that the evidence does not support. A graph also helps follow a chain such as component, batch, test and approved supplier. However, extraction can introduce wrong links, sources can be false and language models can still invent explanations. Graph-based retrieval research supports task-specific improvements; it does not establish that a graph eliminates hallucinations. [8]

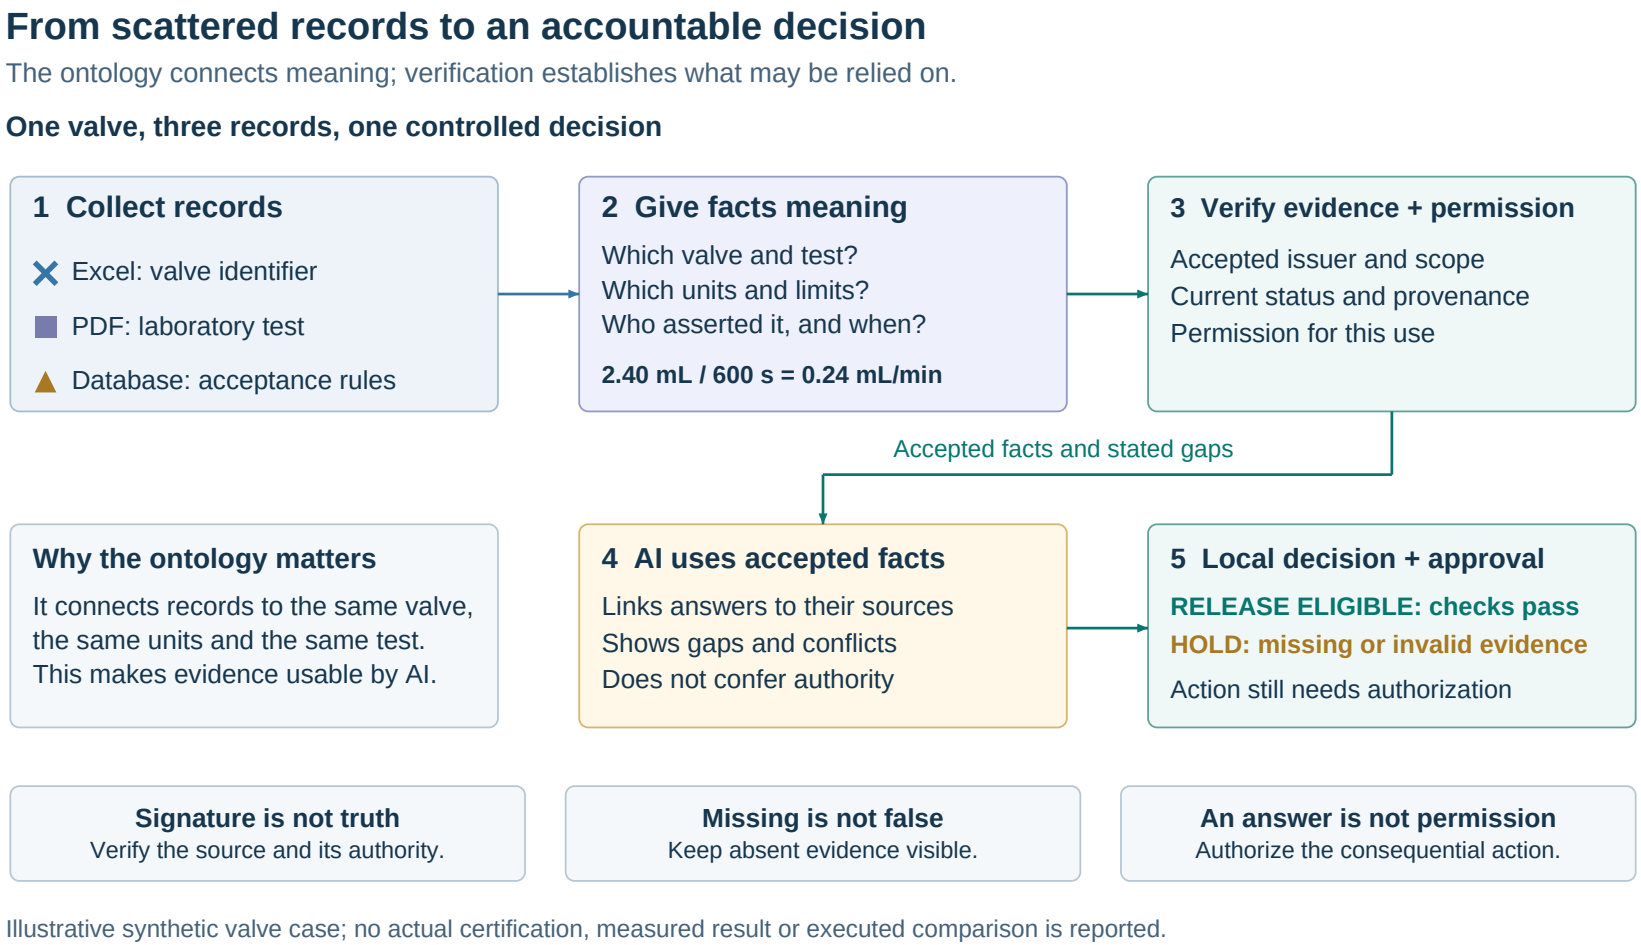


Figure 2: A proposed valve-release example shows how heterogeneous records become decision-relevant evidence. The graph links the subject, measurement, source and applicable rule. Verification and policy checks determine whether the evidence may support the task; AI explains the accepted facts, while the responsible participant controls the operation. A valid signature does not establish truth. Missing evidence can require HOLD. This is an illustrative workflow, not a validated safety case or a reported industrial deployment.

Repeatable checks can make outcomes more robust. Identity matching, unit conversion, freshness checks and acceptance rules can run in explicit software with recorded inputs. This gives reviewers something concrete to inspect and rerun. It is especially valuable when an incorrect answer could affect product release, maintenance, regulatory reporting or a sensitive operating process. The benefit comes from the complete controlled workflow, not the graph format alone.

Faster repeated decisions are possible, but latency must be measured. Reusing an index and an accepted evidence package can avoid repeated document reading. LightRAG reports faster queries than the heavier graph-summarization implementation it tested. More complex retrieval, cross-company queries and verification can also add delay. HippoRAG 2 illustrates the trade-off: stronger results on selected question-answering tasks coexist with greater query time and memory use than a strong dense-retrieval baseline. Procurement should compare end-to-end performance on the actual industrial task. [9] [10]

Risk becomes more visible and manageable. The system can show which claims support a decision, who supplied them and what changes would invalidate them. Missing permission, conflicting evidence or a withdrawn credential can produce HOLD and a human review. These controls can reduce exposure to specific failures. They do not constitute a safety certificate or proof of legal compliance.

The practical performance question is therefore: does this system produce more correct, supported decisions at an acceptable total cost? Measure unsupported material claims, wrong approvals, missed problems, prompt and output tokens, and total operating cost under the same data and access conditions. Report typical response time and the time within which 95% of requests finish. Include failures and denials. A fast answer that uses information the requester was not entitled to see is not a successful result.

The cited studies test retrieval designs. They do not establish that open-world semantics, international standards or cryptographic verification themselves produce the reported AI gains. Those are additional commitments of the proposed ecosystem, with their own tests.

4 A humanoid robot makes the evidence relationships visible

Imagine a fictional industrial humanoid H-01 picking parts. Its component supplier holds the firmware and software inventory; a laboratory holds a test report; the manufacturer assesses the configured robot; the operator records installation and maintenance. Figure 3 connects their selected records. No real robot test, issued certificate or verified signature is presented. Spherity's PACE research supplies the configuration-specific assurance approach; its existing robot-passport viewer is also a concept demonstration. [11] [12]

A report, a claim and an approval are different things. Report TR42 is evidence. Assertion TA42 states what the test found about configuration R3. Release claim RC3 is a proposed conclusion supported by accepted premises under rule AR1. The AI evidence assistant helps follow those relationships; it cannot make a certificate apply to an untested configuration.

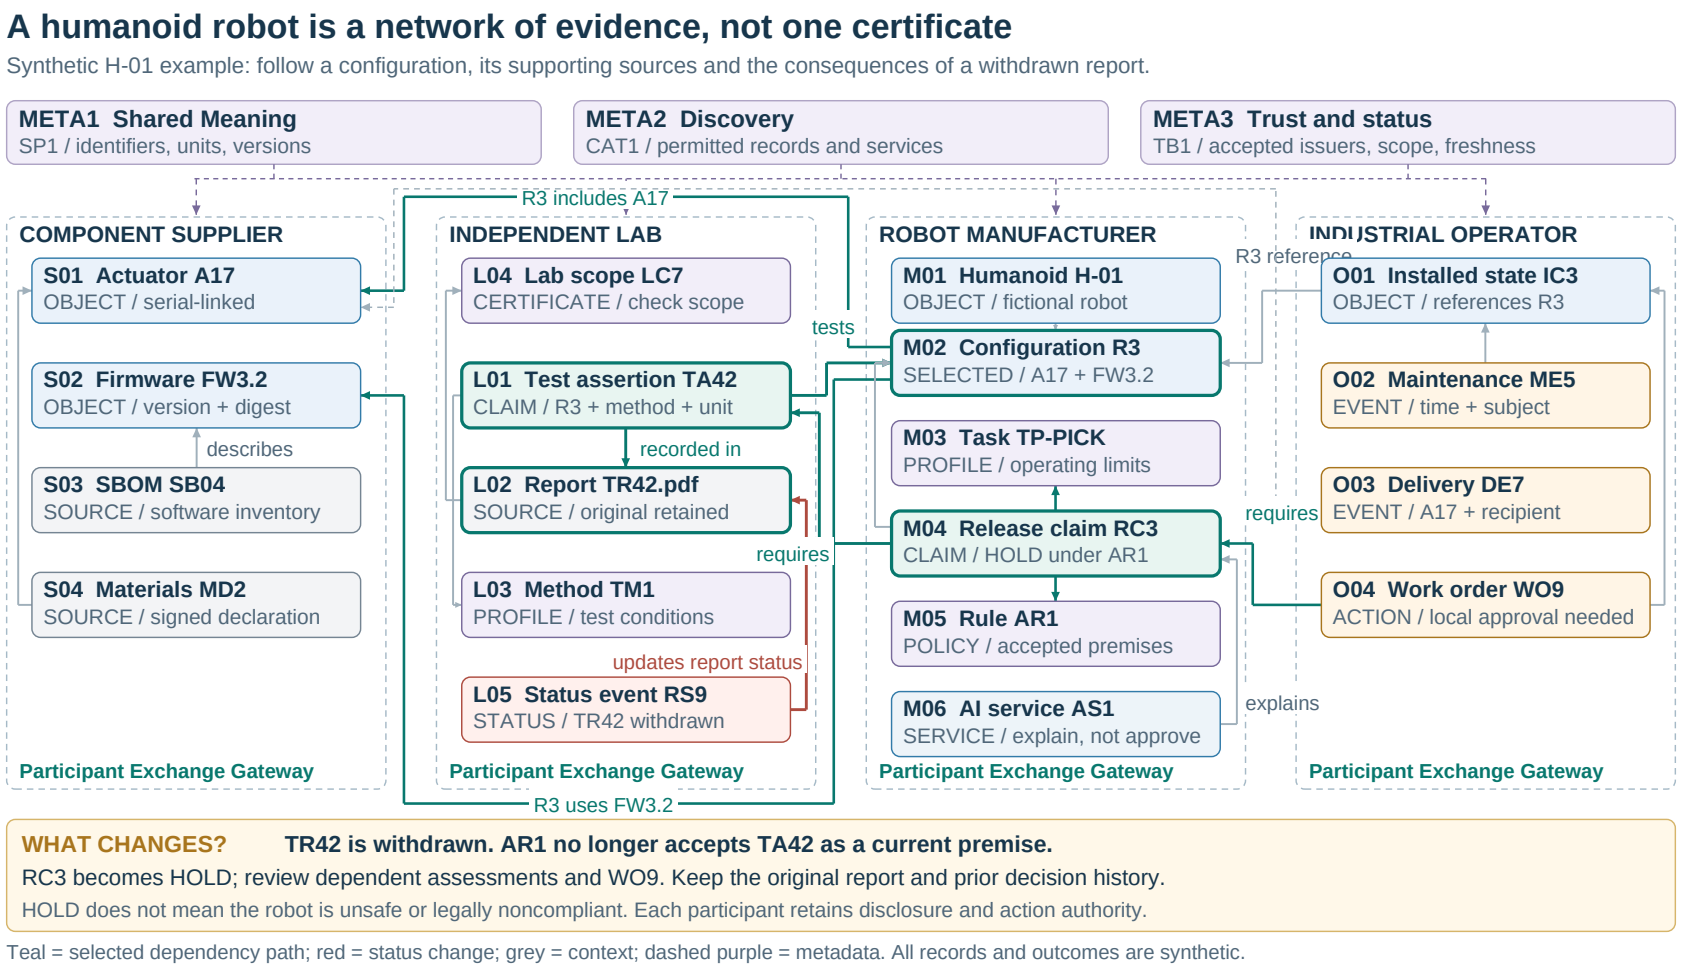


Figure 3: A synthetic humanoid evidence graph connects configuration R3 to components, a test assertion, its original report, scoped authority and operational events across four independently governed participants. Shared semantic, discovery and trust roles supply authorised metadata to all participants; 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.

If the laboratory withdraws the report, the graph shows which decisions depended on it. This example produces HOLD pending reassessment because AR1 requires current test support. It preserves the historical report and unrelated evidence. Withdrawal does not, by itself, prove the robot unsafe.

Figure 4 shows how evidence can serve seven separate assessments: ESPR/DPP, functional safety, the EU AI Act, MDR, machinery, CRA and export controls. **Reuse does not mean one universal approval.** Product scope and intended use still govern: a humanoid has no automatic ESPR passport or high-risk AI classification; current machinery-sector rules matter. [13] [14] [15] A medical variant requires a separate assessment, and CRA excludes MDR-covered products. [16] [17] Functional safety and export screening retain their own criteria. [18] [19]

Each company controls disclosure. AI can use an authorized view to find missing support, explain a change and prepare a dossier, while responsible people retain approval authority and local safety controls remain independent. Test whether reusable evidence makes compliance preparation more automated, lowers cost and time to market, improves risk and cybersecurity management and resilience, and enables product and service innovation. It depends on accurate links and current evidence, and does not mean every AI application requires a knowledge graph.

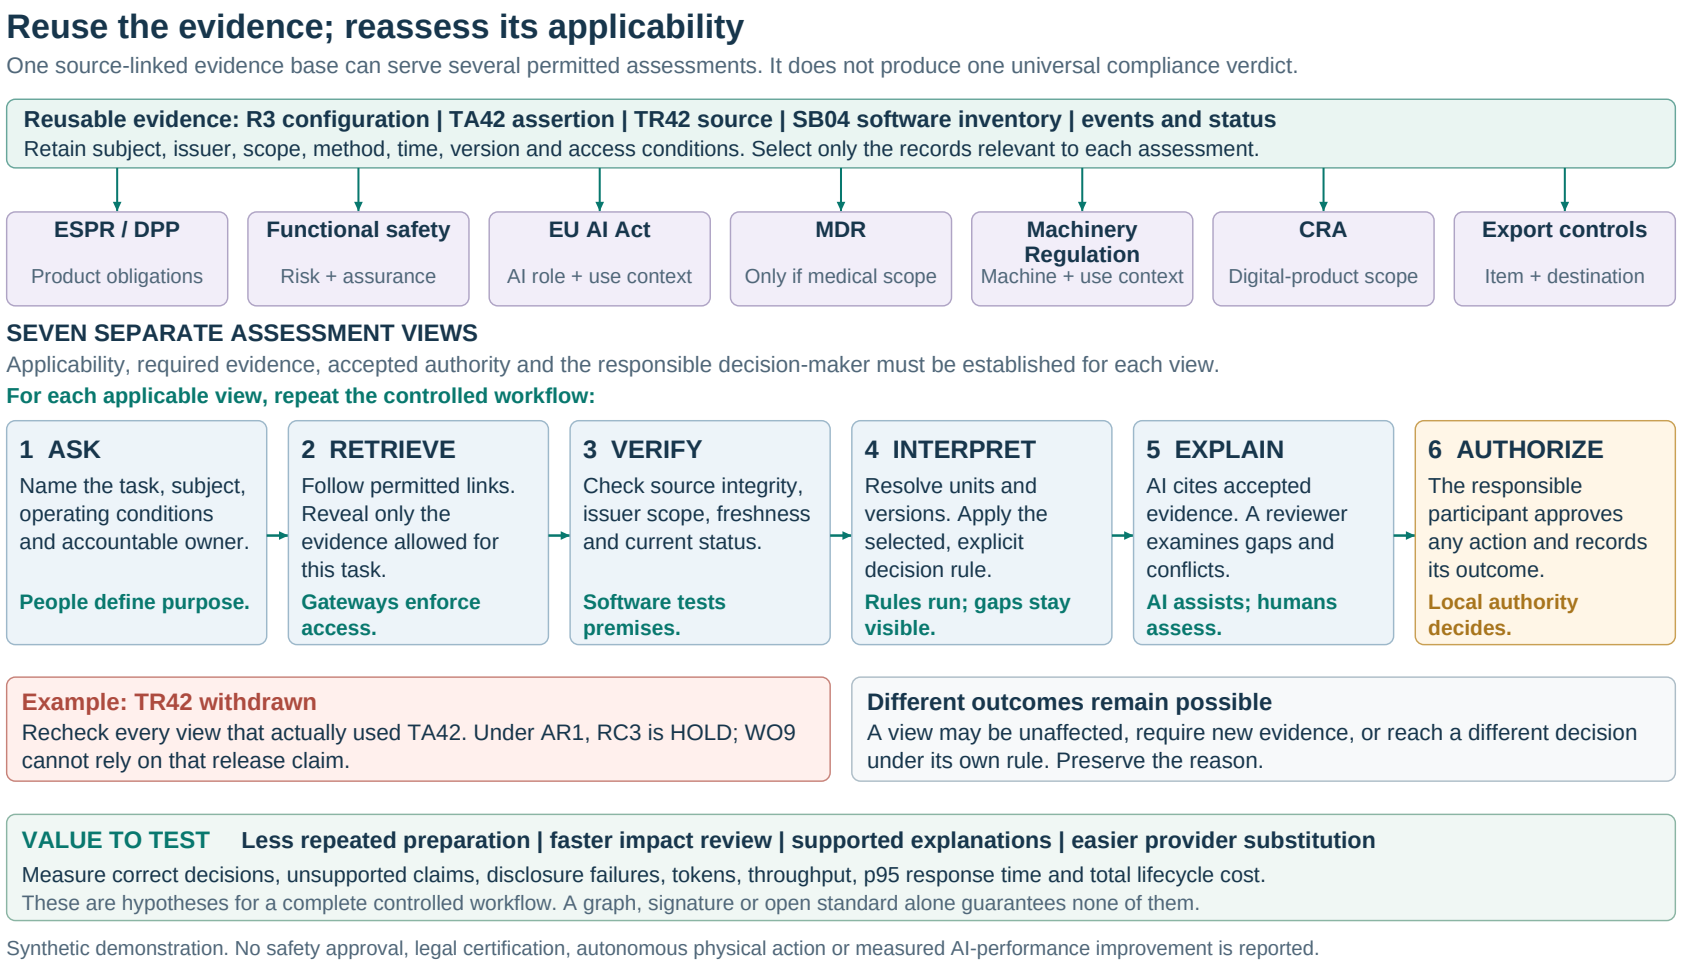


Figure 4: Selected evidence can be reused across seven separate assessment views: ESPR/DPP, functional safety, the EU AI Act, MDR where medical scope applies, the Machinery Regulation, CRA and export controls. Each view requires its own applicability analysis, accepted evidence, authority and decision 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.

5 An ecosystem of ontology engines

The next step is to make this operational capability work across companies. A dataspace organizes data sharing among participants under agreed rules. The proposed architecture builds on that idea: a supplier, manufacturer and laboratory each prepare and link their own records within their governance boundary. They retain their own data, ontology engine and decision authority and may choose different technologies. Local control does not require on-premises hosting. After controlled exchange, an authorized requester can compose a graph view from accepted assertions or query results. Identity matching, compatible vocabulary versions and unresolved conflicts still require explicit checks. No operator needs to possess the full supply-chain graph.

All participants need access to three shared functions. **Shared Meaning** distributes agreed definitions and mappings. **Discovery** helps them locate permitted services and evidence. **Trust and status** help them assess identity, authority and whether information is current. These are shared responsibilities that can have several operators. They need not become one mandatory database or one central gatekeeper.

A **Participant Exchange Gateway** handles the boundary: who is requesting what, on whose behalf, under which conditions and with what evidence. A dataspace connector built with Eclipse Dataspace Components is one implementation option for the exchange functions. Identity services, evidence verification and authorization for consequential actions must be composed with it. A connector alone does not supply the whole architecture. [20]

Business wallets and verifiable credentials can make an organization's identity and scoped authority portable. A verifier must still decide whether the issuer is accepted for the particular claim and purpose. A laboratory's identity does not automatically qualify every test it performs. Standards from the World Wide Web Consortium (W3C) provide international building blocks for verifiable claims; acceptance policies and implementation profiles remain necessary. [21]

Open-world semantics helps express a basic industrial reality: information may exist elsewhere or be withheld. An unavailable test is not automatically a failed test or proof that no test exists. An operational rule can nevertheless require the accepted evidence before proceeding. Open semantics does not mean public data, unrestricted graph traversal or permission to disregard confidentiality. [22]

The shared services also need protection. An altered vocabulary or false directory entry can misdirect an AI workflow even when the underlying business records are secure. Participants need authenticated publishers, checked versions and bounded freshness for local caches. Access should follow the task and the actor's authority. A company should reveal what the task permits, including an appropriately scoped conclusion where full disclosure is unnecessary.

The goal is **reliable reuse across boundaries**. A test result should remain interpretable when a customer changes providers. A supplier should serve several manufacturers without rebuilding every relationship from the beginning. A graph assembled for a permitted task should retain the origin and conditions of its inputs. This is the operational meaning of a globally collaborative counter-model.

6 Fast-track modernisation for industrial AI

Dataspace research was already established in 2005, when Franklin, Halevy and Maier distinguished administrative coordination from semantic integration. These conceptual roots predate

the current frontier-AI and large-language-model wave, not AI itself. Today's European programmes have developed further. Their historical foundations neither establish obsolescence nor demonstrate readiness for demanding AI workloads. [23]

Fast-track modernisation should preserve the cooperation model while testing and upgrading its operational machinery. Define the data volumes, simultaneous requests, response deadlines and evidence-freshness limits for each task. Assess the complete path, including extraction, exchange, verification and authorized action. Faster transport alone does not establish a dependable decision service. Existing components can be retained where they meet those requirements.

Our prior *Evidence Graphs for Industrial AI* proposes separating large or high-frequency raw streams from their decision-relevant evidence. Specialist systems can retain telemetry; the graph connects identifiers, derived measurements, time windows, provenance and references to underlying records. This is a design proposal, not a measured performance result. [24]

The engineering objective is to meet performance targets and trust requirements together. Caching, parallel processing and selective retrieval need explicit limits on freshness and disclosure. Publish the operating conditions, costs and failure behaviour. Coordinate compatible profiles and tests internationally so that modernisation expands practical cooperation across providers and jurisdictions.

7 Digital trust diplomacy in practice

Europe has a policy basis for this offer. The Commission and High Representative's International Digital Strategy of 5 June 2025 connects competitiveness and sovereignty with international partnerships, open standards and cross-border digital public infrastructure. It provides a diplomatic direction, not an endorsement or completed implementation of the architecture proposed here. [25]

Three concrete relationships show how digital trust diplomacy can progress from principles to specific cooperation. The dates and status matter because an agreement to cooperate is different from a working cross-border service.

Japan. On 5 May 2026, the EU–Japan Digital Partnership Council reported a successful pilot exchanging verifiable student academic records through prototype identity wallets. It also identified ongoing cooperation between Japan's Ouranos Ecosystem and Catena-X on battery product carbon-footprint data, and agreed a Data Strategy Working Group. These are concrete steps across independently governed systems. The credential pilot does not establish general production interoperability; the same statement places broader electronic-signature recognition in future discussions. [26]

Singapore. The EU–Singapore Digital Trade Agreement entered into force on 1 February 2026. The Commission identifies paperless trade, electronic transactions and restrictions on unjustified data localisation among its measures. This is an agreement in force, not merely an intention to cooperate. Its legal framework does not itself establish compatible industrial vocabularies, interoperable wallets or acceptance of every digitally signed claim. [27]

India. On 27 January 2026, the Commission's DG CONNECT and India's electronics and IT ministry signed an arrangement on advanced electronic signatures and seals. It structures cooperation on public-key infrastructure, lists of recognised trust-service providers and signature validation. Linking the trusted lists is an expressed intention. This directly addresses cross-border verification, but the announcement does not establish completed linkage or general legal equivalence to EU qualified trust services. [28]

We use **digital trust diplomacy** to describe this combination of technical coordination, institutional acceptance and legal cooperation. The three must be kept distinct. Two systems may exchange a credential successfully while disagreeing about whether its issuer is authorized or whether its claim has a particular legal effect. A useful diplomatic deliverable is a bounded agreement for a real task, with named authorities, tested profiles and a procedure for disputes and corrections.

International participation should be built into the design. US companies and researchers contribute to W3C and other standards communities. Japan, Korea, Singapore, India and China bring industrial partners and their own data-governance priorities. Cooperation with Chinese participants should be scoped to applicable legal requirements and the task's security conditions. A common interface can bridge systems without requiring identical domestic policy or political alignment.

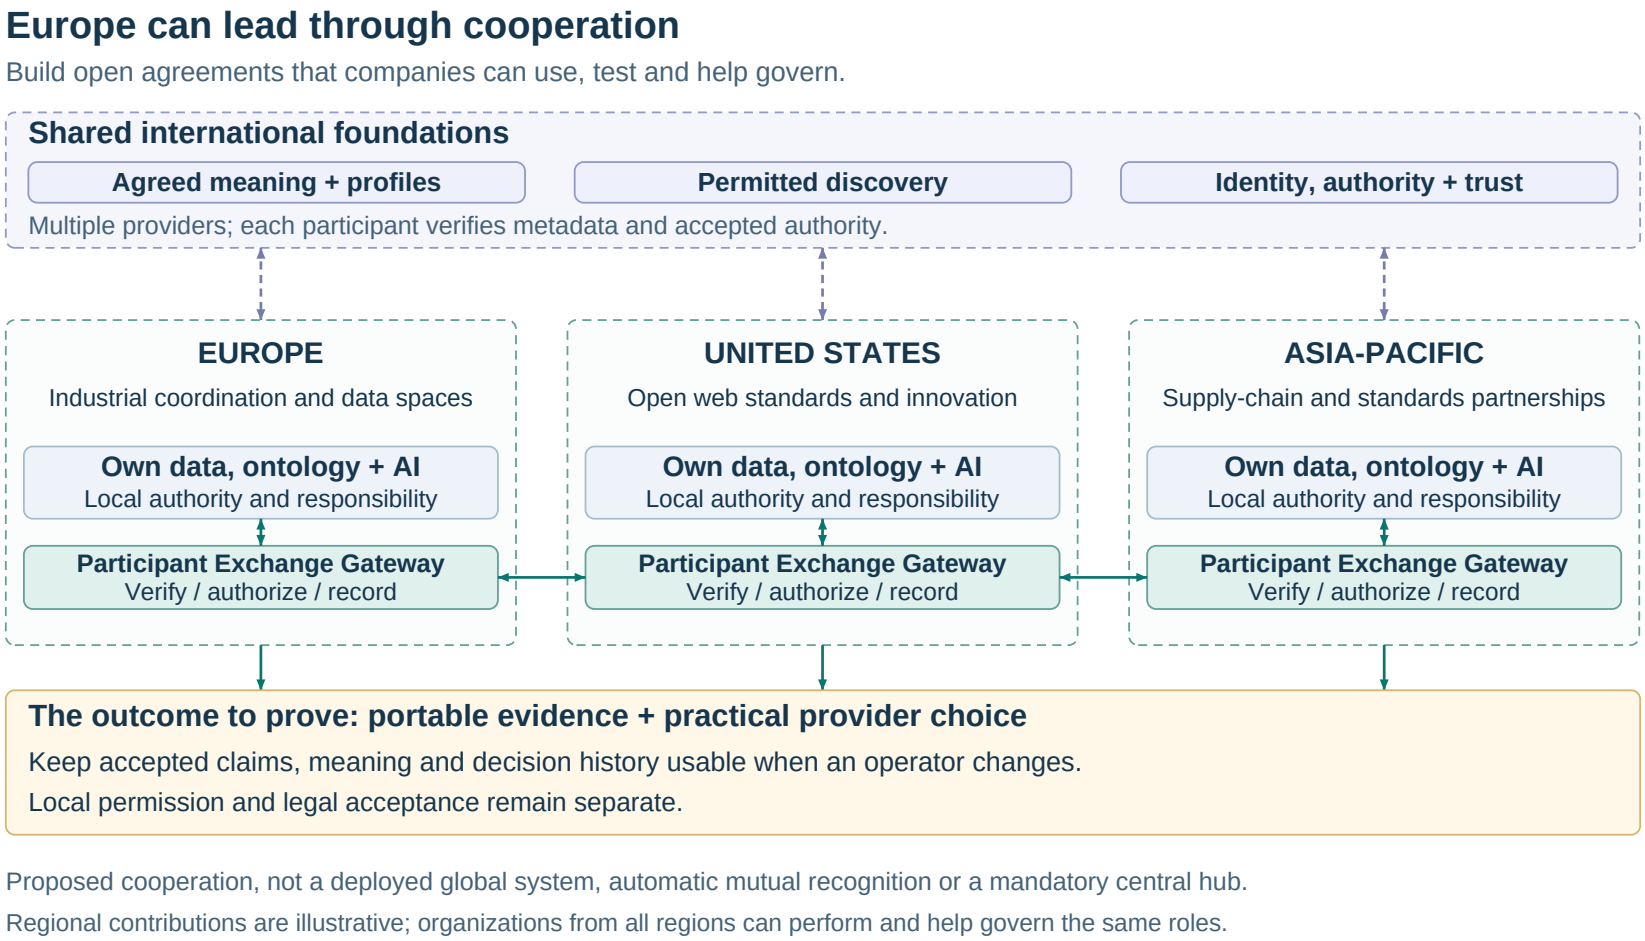


Figure 5: Europe can turn its dataspace experience into a globally collaborative offer. International partners contribute to shared standards and jointly governed interoperability profiles while participants retain local authority. The illustration is a proposed cooperation model. Diplomatic agreements, technical interoperability, accepted issuer authority and legal recognition are separate achievements; their completion is not implied by the connecting lines.

Europe's leadership should be judged by the quality of that bridge. Figure 5 illustrates joint cooperation in which partners help govern the vocabulary, test implementations independently and retain their own decision authority. Publish the supported profiles and unresolved differences. This offers a practical route for supply chains that cross both company and jurisdictional boundaries, while keeping sensitive information under defined control.

8 A different business model needs collaborative leadership

The 2019 paper *Decentralized Identity as a Meta-platform: How Cooperation Beats Aggregation* provides a useful starting point. Shea, Smith and Stöcker propose participant-controlled transfer of value across platforms, using portable identity as a connecting mechanism. Their argument is that cooperation can create useful relationships without requiring one platform to absorb the others. The paper also recognizes costs and difficulties, including onboarding, portability and the user experience. Its network-value argument is conceptual; it is not proof that every decentralized system is cheaper or more valuable. [29]

Our extension is to make **meaning, evidence and decision history** portable alongside identity. A laboratory can supply a reusable, verifiable assertion to several customers. A manufacturer can select a specialist AI service without surrendering its acceptance rules. A software provider can compete by performing a task better, while agreed interfaces allow the customer to substitute it. These are proposed commercial opportunities, not benefits established by membership in a consortium.

This is explicitly a counter-model to Palantir's platform-centred integration and to closed platform ecosystems more broadly. It aims to distribute the ability to create and retain value among participants. Integrated providers, including Palantir, can compete within such an ecosystem where they meet its exchange and portability requirements. The strategic test is whether participants have usable alternatives, including under stress.

Policymakers should procure cooperation and exit as capabilities. Funding and procurement can require a second implementation, portable evidence histories and a demonstrated provider change. Public support should maintain vocabulary stewardship, open test suites and affordable onboarding for smaller firms. A licence label or membership count is insufficient evidence of interoperability. Rewards should follow working exchanges and preserved controls.

Industry managers should fund shared meaning as an operational asset. Domain experts must agree which component a record identifies, what a test measures and how a rule applies. Give vocabulary and mapping owners time and a maintenance budget. Share the non-differentiating definitions and tests; compete on products, services and execution. Where one participant bears integration costs while another receives most of the benefit, agree a workable allocation before scaling.

Europe's IPCEI-AI process can help align investment with this modernisation. The Commission's September 2026 announcement described joint design and planned pre-notification, including decentralized, commonly accessible AI services as an objective. That is an investment direction, not evidence that existing dataspaces, the proposed ecosystem or a particular semantic standard meet the programme's future workload requirements. [30] We recommend funding complete industrial workflows, upgrades where required and independently tested exchange contracts. Demonstrations should meet declared throughput, latency and data-volume targets while preserving access control, verifiable evidence and participant authority; participation in a programme is not proof of readiness.

Collaborative leadership is essential because this technology creates value between organizations. No single firm can unilaterally define a trusted vocabulary for every counterparty. Governance needs a way to add terms, resolve disputes, correct evidence and retire obsolete versions. Implementations of shared technical services should be replaceable while preserving required authority and continuity. Otherwise, a nominally open ecosystem can reproduce the same concentration of power it was meant to address.

9 A practical first mover programme

The following is a proposed 90-day demonstration programme. Its purpose is to establish a bounded working relationship and expose the costs of interoperability and modernisation. Its results would apply to the tested workload, not establish readiness for all industrial AI or IPCEI-AI use cases. Deployment in a safety-critical or regulated process requires the additional assurance and approvals appropriate to that process.

Days 1 to 15 choose one decision and its accountable owner. Bring together a manufacturer or infrastructure operator, a supplier, a laboratory or assurance provider, and at least two technology providers. Select a narrow task such as checking a replacement component's evidence package before a human-authorized release. Record the current cost, error modes and time required. Define realistic data volumes, concurrent requests, response deadlines and evidence-freshness limits. Agree what information may be shared and what must remain private.

Days 16 to 30 agree the minimum common meaning. Define identifiers, units, test scope, the small vocabulary required for the task and the accepted issuers. Reuse relevant standards and industry profiles. Assign a named steward and a version policy. Publish a compact example and conformance tests. The objective is enough agreement to perform the task correctly, not a universal ontology of industry.

Days 31 to 60 connect the participants and keep the evidence. Ingest the same mixed-format records used in normal work. Check extraction and mapping. Configure scoped identity, gateways and access decisions. Retain original assertions, timestamps and transformation history; keep bulk streams in suitable specialist systems where appropriate. Let the AI query and explain accepted facts; use explicit rules and responsible approval for consequential operations. Show an authorized result and a correctly denied request at the declared load.

Days 61 to 75 test the difficult cases. Introduce a wrong identifier, incompatible unit, outdated report, revoked authority and conflicting evidence. Check that missing evidence leads to the agreed HOLD or escalation. Attempt to retrieve information outside the requester's entitlement. Test request bursts, a service outage and a stale trust cache. Record timeouts, denials and whether required evidence and authorization checks remain effective under load. Record failures as well as successes, with no confidential records exposed in public demonstrations.

Days 76 to 90 replace a provider and publish the result. Repeat the agreed decision using a second gateway, graph or AI service at a defined substitution boundary. Preserve the permitted evidence and decision history. Compare correctness, unsupported claims, disclosure failures, tokens, total cost and onboarding effort with the original workflow under matched conditions. Report completed authorized requests per second, typical response time, the time within which 95% finish, and failures separately. Publish data volumes, load conditions, verification costs, the reusable profile and remaining gaps. Expand only when the result meets the agreed performance and trust requirements together.

Beyond the first 90 days, the group needs an operating agreement: maintenance funding, responsibility for incidents, correction and appeal procedures, and a route for new participants

to join. Run at least one cross-border test with independently chosen providers. Track whether the next supplier can join with less repeated work. That is more persuasive evidence of an ecosystem benefit than a larger diagram or a longer list of members.

10 An invitation to build the next operational layer together

Europe's response to Palantir should combine technical ambition with practical choice. Dataspaces are a first step; **fast-track modernisation must turn the cooperation foundation into tested infrastructure for industrial AI**. The aim is high-throughput, responsive operation across company boundaries through shared meaning, verifiable evidence and permissioned action. Access control and participant authority must remain effective at the required scale. US innovators and partners across Asia-Pacific should help shape the standards, modernise implementations and test the resulting services.

We invite first-mover companies and public partners to build and test this together. Bring one consequential use case, a small permitted evidence package and a domain expert who can explain what a correct decision requires. Technology providers should bring an implementation and a willingness to demonstrate substitution. Standards bodies and industry associations should help turn successful profiles into maintained common assets. Policymakers can support neutral testing, smaller-firm participation and cross-border trust arrangements.

Move quickly on a bounded task, share what others need to interoperate and measure performance and trust together. Publish the trade-offs and remaining limits. Collaborative leadership means giving partners a real role in the definitions, implementation profiles and rules that govern cooperation. That is how a European initiative can become a credible global counter-model to platform dependence, rather than another untested promise of readiness.

11 Scope and disclosure

This executive brief translates the accompanying academic paper 1.0 into policy and management recommendations. It combines primary-source documentation, selected AI research and an author-proposed architecture. The illustrations, modernisation agenda and 90-day programme are proposals; no author-run comparative Palantir benchmark or production federation is reported. Sources were reviewed through 1 November 2026. Dr. Carsten Stöcker is affiliated with Spherity GmbH, a provider of digital identity and trust technology, authored the cited evidence-graph research and co-authored the cited 2019 meta-platform paper.

References

[1] Competition and Markets Authority. CMA announces package of actions on business software and cloud services. Published 2026-03-31. Accessed 1 November 2026.

[2] Competition and Markets Authority. CMA takes first steps to ensure choice in business software. Published 2026-10-16. Accessed 1 November 2026.

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

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

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

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

[7] Darren Edge et al. From Local to Global: A GraphRAG Approach to Query-Focused Summarization. Inspected arXiv v2; first submission 2024-04-24. Dated 2025-02-19. Accessed 1 November 2026.

[8] Xiangrong Zhu, Yuexiang Xie, Yi Liu, Yaliang Li, Wei Hu. Knowledge Graph-Guided Retrieval Augmented Generation. Proceedings month; conference 29 April–4 May 2025. Dated 2025-04. Accessed 1 November 2026. Full text.

[9] Zirui Guo, Lianghao Xia, Yanhua Yu, Tu Ao, Chao Huang. LightRAG: Simple and Fast Retrieval-Augmented Generation. Proceedings date; conference 4–9 November 2025. Dated 2025-11. Accessed 1 November 2026. Full text.

[10] Bernal Jiménez Gutiérrez, Yiheng Shu, Weijian Qi, Sizhe Zhou, Yu Su. From RAG to Memory: Non-Parametric Continual Learning for Large Language Models. Inspected arXiv v2. Dated 2025-06-19. Accessed 1 November 2026.

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

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

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

[14] 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.

[15] 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.

[16] 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.

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

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

[19] 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.

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

[21] W3C. Verifiable Credentials Data Model 2.0. Published 2025-05-15. Accessed 1 November 2026.

[22] W3C. OWL 2 Primer. Published 2012-12-11. Accessed 1 November 2026.

[23] 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.

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

[25] 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.

[26] European Union, Japan. Joint Statement of the fourth meeting of the European Union–Japan Digital Partnership Council. Published 2026-05-05. Accessed 1 November 2026.

[27] European Commission, Directorate-General for Trade and Economic Security. EU–Singapore Digital Trade Agreement enters into force. Published 2026-02-02. Accessed 1 November 2026.

[28] European Commission, Digital Building Blocks. EU and India Advance Cross-Border Digital Trust Through an Administrative Arrangement on Advanced Electronic Signatures and Seals. Undated / rolling source. Accessed 1 November 2026. Updated 2026-02-06.

[29] 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.

[30] 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.