

Can innovation commons be managed digitally – and should they be?

Thomas Wernerheim, Lars Boström, Jakob Lindvall



compare

DIGITALWELL
ARENA

IDEON
SCIENCE PARK

VINNOVA

1. For the chronically impatient

The project Development and Application of Innovation Commons for Strengthening Innovation Eco-Systems was conducted by Compare/DigitalWell Arena, Karlstad Municipality, and Ideon Science Park AB from December 2024 to December 2025, funded by Vinnova. The project was structured around three interrelated work packages. **WP1** focused on sharing, developing, and disseminating insights on commons for more efficient innovation ecosystems. **WP2** focused on the development of innovation commons management. **WP3**, which this report documents, focused on piloting a technical infrastructure for managing commons in innovation ecosystems.



This report examines whether digital infrastructure can support innovation commons systems, and under what institutional conditions. Rather than presenting a finished solution, it documents what became visible when WP3 combined structured market dialogue, hands-on prototyping, and use in a live community context.

Three components were explored. The **Knowledge Manager** was tested as an institutional prototype for AI-supported retrieval and synthesis over curated artefacts. The **community platform** was explored as an enabling space for participation, facilitation, and stewardship. The analysis of **distributed knowledge nodes and domain-level exchange** added a system-level perspective on how value may be formed locally while still remaining shareable across domains. In this report, the analysis primarily operates at the level of artefacts and

repositories, since this is the level at which the technical infrastructure actually manages resources through documents, metadata, traceability, and licensing conditions.

Taken together, the findings suggest that **meaningful value does not emerge from technology alone**. AI-supported knowledge services can create practical value when grounded in curated artefacts, stable definitions, traceability, and metadata discipline. Community platforms can support participation, visibility, and stewardship, but do not in themselves generate sustained collective action. Distributed nodes may preserve local context and autonomy, yet require interoperability, licensing coherence, and governance alignment if exchange is to work across the wider innovation commons system.

A common pattern runs through the report: **the decisive constraint is not technical feasibility as such, but institutional maturity**. Governance arrangements, stewardship capacity, participation logic, and disciplined resource management shape whether digital infrastructure strengthens pooling and reuse over time or merely adds another technical layer. At the same time, the report indicates that some governance questions extend beyond artefact-level infrastructure to the level of informational content, especially where AI-supported synthesis cuts across multiple sources.

The report also clarifies what remains unresolved. Conversational material may represent an important extension of the innovation commons resource, but raises more complex questions regarding licensing conditions, traceability, and stewardship. Cross-node exchange appears promising, but also increases governance complexity. For these reasons, the report should be read as an institutional learning process and a basis for continued hypothesis testing rather than as a blueprint for implementation.

Content

1.	For the chronically impatient	2
2.	A Question, Not a Summary	5
2.1	<i>What We Were Trying to Achieve (and Why It's Hard)</i>	6
2.2	<i>How We Chose to Explore the Problem</i>	7
3.	PART I – Exploring the Knowledge Manager	8
3.1	<i>The Knowledge Problem We Tried to Solve</i>	8
3.2	<i>What the Market Told Us (and What It Didn't)</i>	8
3.2.1	Incentives and Participation Logic	9
3.2.2	Governance Without Centralisation	10
3.2.3	Architecture as Operationalised Governance.....	10
3.2.4	Incentives and Actor Motivations: Why Participate? (Redpill Linpro)	11
3.2.5	Federated Governance and Trust Infrastructure (AI-Bees)	13
3.2.6	Technical Realisation and Digital Sovereignty (Sweco)	14
3.2.7	Distributed Knowledge Nodes and Domain-Level Exchange	16
3.2.8	Concluding Reflection: What the Dialogue Added	18
3.2.9	Synthesis	19
3.3	<i>Building and Testing the Knowledge Manager.....</i>	19
3.3.1	Purpose and scope of the Knowledge Manager	20
3.3.2	Technical approach and architecture	21
3.3.3	Testing, evaluation, and technical limitations	23
3.3.4	Licensing, conditions, and knowledge commons	27
3.3.5	Relation to the community platform.....	27
3.3.6	Summary of key learnings	28
4.	PART II – Exploring the Community Platform	30
4.1	<i>Why Technology Alone Was Never Enough</i>	30
4.2	<i>What Happened When People Actually Used It.....</i>	31
4.3	<i>What We Learned About Communities and Commons</i>	32
4.4	<i>Research context: participation and value creation in community platform</i>	34
4.5	<i>The Platform as a Monitoring and Evaluation Infrastructure</i>	35
5.	PART III – When AI Meets Community	37
5.1	<i>The Interplay: Where the Real Value Emerges</i>	37
5.2	<i>What This Means for Governance and Sustainability</i>	38
5.3	<i>What We Would Recommend (and What We Would Not)</i>	39
6.	Closing Reflections: From Pilot to Possibility.....	40
	References	41

2. A Question, Not a Summary

Can innovation commons be managed digitally, and should they be?

This report opens with a question rather than a summary because the question is genuine. When WP3 set out to explore whether digital infrastructure could support innovation commons systems, the answer was not obvious. It still is not.

The final report of the project, *Development and Application of Innovation Commons for Strengthening Innovation Eco-Systems*, establishes a foundational premise: an innovation commons is not simply a shared resource. It is a governance system in which a community pools distributed information under shared rules to enable entrepreneurial discovery that no single actor could generate alone. Technology may serve such a system. It may also undermine it. Which outcome occurs depends not on the technology itself, but on the institutional conditions under which it is deployed.

WP3 was the part of the project that put this directly to the test. Two artefacts were explored: a **Knowledge Manager**, designed to make curated resources navigable through AI-supported retrieval, and a **Community Platform**, designed to enable sustained participation and stewardship. Neither was conceived as a finished product. Both were designed as learning experiments. This report documents what became visible when we tried.

In this report, the analysis **primarily operates at the level of artefacts and repositories**, since this is the level at which the technical infrastructure actually manages resources through documents, metadata, traceability, and licensing conditions. **The broader governance framework developed in the final report extends this perspective by analysing governance also at the level of information objects**, that is, the informational content that artefacts contain and circulate.

Three tensions that will not go away

Knowledge, in an innovation commons system, is abundant. What it is not is usable. Accumulated artefacts exist in growing quantities yet remain hard to navigate and rarely reused in ways proportionate to the effort that produced them. The problem is not knowledge scarcity. It is institutional incapacity to activate distributed information over time.

Communities exist, but are not self-sustaining. They require facilitation, stewardship, and ongoing governance to remain viable as commons rather than becoming spaces where information is consumed without being replenished.

Technology promises scale, but risks oversimplification. AI-supported retrieval and modular architecture hold genuine potential, yet scalability without appropriate governance arrangements may expand access while eroding the stewardship conditions that make a commons viable.

These three tensions were identified in the pre-study on information structures in learning networks (Wernerheim et al., 2024), which preceded this project and shaped its design. WP3

built on that foundation by moving from conceptual analysis to applied experimentation within a live innovation commons system.

The hypothesis this report tests

The final report frames the project around a central question: what does it take to cultivate innovation commons deliberately, rather than waiting for them to emerge spontaneously? Jason Potts observed that such commons tend to arise organically in the zeroth phase of new technological fields but had not been seen sustained intentionally over time. WP3 carries a specific version of this question into the technical domain.

The hypothesis is narrow and testable: under the right institutional conditions, digital infrastructure may reinforce the pooling and reuse of distributed information in ways that strengthen rather than substitute collective action. The evaluation framework developed for the project, grounded in learning-oriented evaluation rather than compliance measurement, oriented the inquiry accordingly. The question is not whether the technology worked in a technical sense, but what the process of trying revealed about governance conditions and sustainability.

2.1 What We Were Trying to Achieve (and Why It's Hard)

Most existing platforms were not designed for commons governance. A document repository treats knowledge as a stock to be stored and retrieved. A social network treats participation as engagement to be maximised. A commons treats distributed information as a resource requiring ongoing stewardship, clear rules of access and attribution, and governance mechanisms that can adapt as the community evolves. The design challenge is therefore not primarily technical. It is institutional.

The theoretical grounding for this claim runs through two bodies of work. Ostrom's design principles for robust commons, developed for natural resource governance, carry clear implications for digital environments where knowledge, data, and artefacts constitute shared resources. Principles such as clearly defined boundaries, participation in rule-making, monitoring, and conflict-resolution mechanisms point not only to the need for social norms, but to explicit institutional arrangements that must be reflected in platform architecture, roles, and governance frameworks (Ostrom, 1990; Ostrom, 2005). Within the knowledge commons literature, Hess and Ostrom further emphasise that digital artefacts, despite being non-rivalrous and easy to copy, require clearly articulated rules for access, use, attribution, and reuse in order to avoid quality degradation, free-riding, and coordination problems (Hess & Ostrom, 2007).

Licensing and rights structures are one such institutional mechanism. They operationalise sharing rules in practice and are critical for enabling collaboration, quality assurance, and scalable cooperation in digital communities (Madison, Frischmann & Strandburg, 2014). This implies that license management is not a peripheral legal add-on but an integral part of the community's institutional rule system. Platform architecture must therefore support a clear linkage between roles, rights, and licensing conditions through features for attribution, version control, access management, and transparency regarding how artefacts may be used and further developed.

Three structural tensions follow from this institutional framing. Openness versus quality: a commons benefits from wide contribution, but openness without curation risks eroding the definitional stability that makes a shared knowledge base genuinely useful. This cannot be resolved by better search algorithms. It requires governance. Automation versus stewardship: AI-supported retrieval can lower participation barriers, but may also produce transactional patterns of engagement that extract from the commons without replenishing it. Stewardship requires human judgement and cannot be automated away. Scalability versus contextual relevance: scale that is not grounded in shared institutional context risks producing a commons without community, a platform that aggregates information without generating the collective intelligence that pooling under shared governance is supposed to enable.

WP3's role within the project was to test these tensions in practice. The final report develops the governance framework. WP3 asked what happened when we actually tried to build something, and what the attempt revealed that theory alone could not.

2.2 How We Chose to Explore the Problem

We did not design a system. We designed learning experiments.

An experimental approach was chosen because the problem is complex in a specific sense: the relevant variables cannot be controlled, causal chains are not linear, and findings will be contextually conditioned. The exploration proceeded through three mechanisms.

- Hands-on technical experiments: the team built, configured, and used the artefacts within an active community context, producing insights about governance-technology interaction that an evaluation exercise would not have surfaced.
- Structured dialogue with the market through a Request for Information (RFI), which functioned as an entrepreneurial discovery mechanism rather than a procurement step: a way to test the project's problem framing against external expertise and surface feasible governance arrangements and architectural patterns.
- Real use in a live community context, primarily the Demand Acceleration Community, grounding the experiments in actual participation patterns and stewardship demands rather than hypothetical use cases.

WP3's work was embedded within the project's broader architecture. WP1 developed the theoretical and governance framework that oriented WP3's design choices. WP2 provided the community context in which governance assumptions could be tested. The cross-workpackage interaction is itself an instance of the pooling under shared governance that the project studies.

What this approach made visible, and what a linear implementation project would not have, was the gap between how governance is intended and how it functions in use. These are signals from practice, exactly the kind of adaptive information that a learning-oriented evaluation is designed to surface. What follows is not a description of what was built, but an account of what the building revealed.

3. PART I – Exploring the Knowledge Manager

3.1 The Knowledge Problem We Tried to Solve

Innovation commons depend on distributed information becoming usable across actors, roles, and contexts. In practice, however, accumulated artefacts often remain fragmented, difficult to navigate, and unevenly reused. The challenge is therefore not absence of knowledge, but limited institutional capacity to activate it over time.

Within the Demand Acceleration context, a growing body of curated documents — guides, case material, and methodological artefacts — had been produced. These artefacts constituted an innovation commons resource, yet access required prior familiarity with terminology, structure, and informal context. New participants faced orientation barriers, while experienced actors relied on tacit understanding. This created structural asymmetries in reuse and participation.

This pattern mirrors a broader challenge identified in publicly funded innovation programmes, including those supported by Vinnova: substantial knowledge production within projects, but limited diffusion, scaling, and long-term institutionalisation of results. Project-based logics tend to prioritise delivery over sustained pooling and reuse.

The earlier pre-study (Wernerheim, et al., 2024) *Information Structures for Learning Networks (ISLNd)* indicated that AI-supported knowledge services require structured metadata, clearly defined responsibilities, and governance alignment in order to generate reliable value. WP3 built on this insight by moving from conceptual analysis to applied experimentation within an existing innovation commons system.

The knowledge problem was therefore framed as institutional rather than purely technical:

Under what conditions can an AI-based Knowledge Manager strengthen participation logic, reinforce innovation commons governance, and support long-term viability of a shared knowledge base — without weakening stewardship or creating false authority?

The Knowledge Manager was thus explored not as an autonomous solution, but as a socio-technical component embedded within an innovation commons system.

3.2 What the Market Told Us (and What It Didn't)

The project reframed the knowledge challenge as an institutional infrastructure question within a socio-technical system rather than as a purely technical build task. This reframing informed the decision to engage the market through a Request for Information (RFI). The RFI did not function as a conventional procurement step. It operated as a structured mechanism for entrepreneurial discovery — a way to test the project's problem framing against external expertise and to surface feasible governance arrangements and architectural patterns.

The RFI responses did more than confirm technical feasibility. They indicated three recurring design layers that future implementation must reconcile:

- incentives and participation logic
- innovation commons governance arrangements that enable trust without centralisation
- technical architecture capable of operationalising rules through traceability and licensing conditions

Across respondents, these layers were emphasised differently. Taken together, they formed a coherent progression from why actors would contribute to an innovation commons resource, to how actors could collaborate without losing control, to how such collaboration could be embedded in code.

At the same time, the market dialogue revealed a limitation. Several proposals implicitly assumed that once the correct institutional infrastructure and technical structure are in place, collective action and sustained pooling will follow. The project treated this assumption as indicative rather than conclusive. Participation logic appeared repeatedly as a design outcome, while the institutional conditions required for long-term viability were less specified.

3.2.1 Incentives and Participation Logic

The first layer concerned participation. Before architecture and before governance detail, the RFI responses returned to a prior question: why would actors contribute to a shared knowledge infrastructure?

Across the material, participation was framed primarily as a value question. Reuse, visibility, recognition, and relevance were identified as central drivers. Contribution became more likely when actors could create operational or reputational value, when quality was safeguarded, and when responsibility structures were clear. This suggests that participation logic is closely linked to actors' ability to capture, claim, and control their contributions.

Three enabling conditions were repeatedly highlighted:

- clear articulation of value creation across actor boundaries
- mechanisms for moderation, review, and quality assurance
- transparent recognition of contributions linked to utility rather than volume

The dialogue also exposed a potential blind spot. Participation was sometimes treated as something that can be activated through interface design or contribution metrics. The findings suggest this is only partially valid. Incentives are embedded in institutional conditions, legitimacy, and stewardship. Participation logic therefore functions as an infrastructural design question rather than a post-deployment adoption issue.

3.2.2 Governance Without Centralisation

The second layer concerned innovation commons governance: how collaboration can scale without requiring actors to surrender autonomy.

Several responses converged around federated models. Rather than centralising all knowledge, respondents proposed architectures based on distributed knowledge pods, policy-driven access management, and role-based responsibility structures. The intention was to enable pooling while preserving institutional control.

This reflects a broader recognition that trust cannot rely solely on informal relationships at scale. Instead, governance arrangements must be embedded into system design through certificates, defined roles, compliance mechanisms, and explicit policy frameworks.

The RFI dialogue also made a structural tension visible:

- stronger governance mechanisms may increase trust
- excessive formalisation may reduce flexibility and entrepreneurial discovery

The implication is that governance arrangements must remain modular and adaptive. Innovation commons systems evolve over time; their institutional infrastructure must allow rules and responsibilities to evolve alongside them.

3.2.3 Architecture as Operationalised Governance

The third layer concerned technical architecture as the operationalisation of governance arrangements.

Respondents emphasised modular design, API-based integration, version control, and structured metadata as foundational components. Architecture was framed not only as infrastructure for scale but as a mechanism for embedding rules directly into system logic, enabling traceability and enforceable licensing conditions.

For WP3, this shifted the focus from model selection toward executable governance. Several principles appeared consistently across responses:

- modular components to reduce dependency and avoid lock-in
- separation between content, metadata, and access logic
- explicit linkage between artefacts, roles, and licensing conditions
- traceability through versioning and container logic

The responses indicate that building an AI-enabled Knowledge Manager is technically feasible. What remains unresolved is how architectural robustness translates into sustained collective action. Architecture can enforce rules; it cannot in itself generate legitimacy or shared purpose.

3.2.4 Incentives and Actor Motivations: Why Participate? (Redpill Linpro)

Redpill Linpro's contribution examined the practical conditions under which a shared digital knowledge environment becomes meaningful for potential contributors. The analysis consistently treats the Knowledge Manager not as a standalone tool but as part of a broader socio-technical system in which participation logic, governance arrangements, and institutional credibility must co-evolve.

A central observation concerns the relationship between user heterogeneity and information structure. The report indicates that a broad and cross-sectoral target group significantly increases the requirements for how distributed information is organised and exposed. Users must be able to orient themselves quickly and with high precision; otherwise, the platform risks being perceived as an additional information layer rather than an enabling innovation commons resource.

To address this, Redpill Linpro emphasises the importance of a disciplined information architecture built on:

- consistent metadata models
- machine-readable tagging
- open indexing that supports context-sensitive search

(Redpill Linpro, 2023)

The implication is that semantic interoperability functions as a foundational institutional condition for scalable pooling and reuse. Without shared conceptual structures and metadata discipline, AI-supported retrieval is unlikely to deliver sustained value. The report therefore positions metadata quality not as a technical detail but as part of the institutional infrastructure required for long-term viability.

In its value analysis, Redpill Linpro draws a clear distinction between individual and organisational benefit. At the individual level, value is primarily associated with reduced search time and improved access to quality-assured knowledge artefacts. At the organisational level, the emphasis shifts toward reduced duplication of effort and improved coordination across organisational boundaries.

This dual value logic suggests that participation in the innovation commons community depends on visible benefits at both levels. Contributions must generate local utility while also strengthening the collective knowledge base. The platform's role is therefore less about central accumulation and more about enabling distributed reuse across contexts and lifecycles.

Quality assurance and trust are treated in relatively operational terms. Redpill Linpro proposes a review and moderation model inspired by established open source practices, where responsibility for artefacts is explicitly assigned and changes remain traceable over time. The proposed model includes:

- designated maintainers or information owners
- open review workflows prior to publication
- systematic version control of artefacts
- AI-supported anomaly detection

(Redpill Linpro, 2023)

A notable aspect of the proposal is that quality assurance is framed as a continuous stewardship process rather than a gatekeeping event. AI is positioned as support for detection and prioritisation, while human oversight remains necessary for interpretation and accountability. This aligns with the broader observation that technical automation can strengthen traceability but cannot replace institutional judgement.

Measurement is approached with explicit caution. Redpill Linpro recommend transparency in usage statistics and contribution visibility but warn against equating activity volume with value creation. Recognition mechanisms should instead be linked to demonstrated utility and quality. This reflects an awareness that poorly calibrated metrics may distort participation logic and incentivise low-value contributions.

The report also highlights structural tensions in cross-sector collaboration. Differences between public and private actors are described in terms of transparency obligations, protection of competitive assets, risk tolerance, and planning horizons. These asymmetries are treated as persistent institutional conditions rather than temporary frictions.

Particular attention is given to information classification. Ambiguity regarding what can be shared, reused, or restricted is identified as a primary barrier to pooling. The report therefore emphasises the need for early and explicit definitions of openness levels, licensing conditions, and role-based responsibility. From an innovation commons governance perspective, this clarity functions as a prerequisite for predictable participation.

Finally, Redpill Linpro stress that interoperability must be addressed across multiple layers simultaneously. The report distinguishes between:

- technical interoperability (APIs and formats)
- semantic interoperability (concepts and metadata)
- organisational interoperability (roles, processes, and governance arrangements)

This multi-layer framing is analytically significant. It indicates that sustainable distributed information exchange depends on alignment across the socio-technical system rather than on technical integration alone.

Taken together, Redpill Linpro's analysis clarifies the institutional conditions under which participation in an innovation commons system becomes plausible. The findings suggest that

incentives cannot be engineered solely through interface design or AI capability. Long-term viability instead depends on disciplined metadata practices, visible value creation, and governance arrangements that make contribution both safe and meaningful for participating actors.

3.2.5 Federated Governance and Trust Infrastructure (AI-Bees)

AI-Bees addressed a structural constraint that recurred throughout the RFI dialogue: organisations are expected to participate in pooling and distributed information sharing while simultaneously maintaining control over sensitive data, models, and institutional responsibilities. Their proposal can be interpreted as an attempt to stabilise this tension through a federated innovation commons system.

The architectural core is the concept of AI pods. Each participating organisation operates its own execution and data domain, within which storage, model execution, and access control remain locally governed. External exposure occurs selectively and under explicit policy conditions. This design choice reflects a deliberate prioritisation of institutional autonomy as a precondition for sustained collective action.

In practical terms, the pod model implies that:

- data remains locally stored
- models execute within the local organisational domain
- access policies are defined and enforced locally
- only selected knowledge artefacts are exposed externally

(AI-Bees, 2025)

From an innovation commons governance perspective, the pods function as bounded participation environments. They allow actors to contribute to a shared innovation commons resource without relinquishing ownership or operational control. The implication is that pooling becomes conditional and policy-mediated rather than structurally centralised.

The trust framework constitutes the institutional backbone of the proposal. AI-Bees describe an infrastructure in which identities, roles, and permissions become cryptographically verifiable and machine-readable. The intention is to reduce ambiguity in distributed information exchange and to strengthen traceability across organisational boundaries.

Core elements of the framework include:

- verifiable organisational identities
- certificate-based role assignment
- machine-readable policy rules
- automated validation at the point of knowledge exchange

(AI-Bees, 2025)

This represents a shift toward infrastructure-mediated trust. Verification mechanisms can reduce coordination costs and increase predictability in cross-organisational interaction. At the same time, the material suggests an important boundary condition: technical verification does not in itself generate legitimacy or willingness to participate. Stewardship and institutional alignment remain necessary complements.

AI-Bees explicitly connect the trust framework to innovation commons governance through policy-driven sharing. Shared rules can be defined at ecosystem level while enforcement remains local via validation contracts. This layered model allows common governance arrangements to coexist with organisational autonomy, which is consistent with the federated participation logic observed across the RFI responses.

The report also gives relatively detailed attention to economic sustainability. AI-Bees outline a multi-layer model combining:

- differentiated membership tiers
- platform service fees
- marketplace-based value creation

(AI-Bees, 2025)

The marketplace component is particularly significant. It introduces a mechanism through which participants can both consume and provide specialised capabilities within the same institutional infrastructure. In principle, this may strengthen incentives for contribution and support long-term viability of the innovation commons system.

At the same time, the material remains indicative regarding several operational questions. These include how quality assurance of marketplace contributions will be maintained, how participation logic evolves as commercial incentives grow, and how stewardship responsibilities are distributed across the federation. These uncertainties do not weaken the architectural proposal but suggest areas where empirical validation will be required.

3.2.6 Technical Realisation and Digital Sovereignty (Sweco)

Sweco's contribution provides the most technically detailed analysis in the RFI material and can be read as an effort to make innovation commons governance executable within a production-oriented socio-technical system. Whereas the preceding contributions focus primarily on participation logic and governance design, Sweco concentrates on how these principles can be embedded directly into architectural structure.

A central design principle is modularity. The proposed architecture decomposes the platform into loosely coupled components covering identity management, data handling, AI integration, and deployment pipelines. Each component is intended to be replaceable without requiring systemic redesign (Sweco, 2026). This reflects an explicit attempt to align technical flexibility with long-term institutional viability.

From a sustainability perspective, modularity serves several functions. It reduces vendor dependency, supports incremental evolution, and preserves optionality in a rapidly changing AI landscape. The architecture therefore treats technical change as an expected condition rather than an exception.

The implementation builds on established cloud-native patterns, including containerisation, orchestration, and standards-based identity protocols. Services communicate through explicit APIs rather than shared databases. This separation is analytically important: it mirrors organisational boundaries and supports role-based responsibility across the innovation commons system.

Security by default is operationalised through strict project isolation. Each service — including AI components — operates within its own authorisation domain and runtime boundary. Sweco describe this in terms of limiting blast radius (Sweco, 2026). In multi-actor environments handling distributed information, such containment mechanisms are likely to be a precondition for institutional trust.

Digital sovereignty is treated as an architectural requirement rather than a purely strategic preference. Sweco argue that the platform should be capable of operating fully within EU jurisdiction if required. To support this, the report evaluates several deployment paths:

- EU-based public cloud providers
- EU bare-metal infrastructure
- self-hosted platform stacks
- carefully bounded use of hyperscalers

(Sweco, 2026)

Importantly, sovereignty is framed as a spectrum rather than a binary state. Architectural discipline can progressively reduce dependency even when pragmatic choices initially involve global cloud providers. This staged approach aligns with the project's experimental posture and supports gradual institutional alignment.

One of the more structurally significant proposals is constitution-as-code implemented through GitOps. In this model, the desired state of the platform — including infrastructure configuration and deployment logic — is declared in version-controlled repositories and automatically reconciled into the runtime environment after governed approval.

In practice, this enables:

- reproducible infrastructure states
- strengthened traceability across the lifecycle
- reduced dependence on individual operators
- controlled evolution of governance-related configuration

(Sweco, 2026)

The approach effectively embeds parts of innovation commons governance into executable artefacts. However, the material also indicates an important limitation: automation strengthens consistency but does not replace stewardship. Institutional conditions and participation logic still require active human oversight.

Sweco complement this with a contract-based architecture. OpenAPI specifications and service-level objectives define formal interaction boundaries between components and, by extension, between organisational actors. As long as these contracts are maintained, internal implementations can evolve independently without destabilising the overall system.

From a socio-technical perspective, contract discipline functions as an enabling condition for distributed collaboration. It allows pooling of capabilities while preserving organisational autonomy. At the same time, it introduces an ongoing governance responsibility: contracts must be versioned, monitored, and institutionally maintained over time.

Taken together, Sweco's analysis indicates that the technical foundations for a robust innovation commons system are largely available. Modular cloud-native architecture, traceable pipelines, and policy-aware identity layers can support distributed information at scale. What remains open is how these technical capabilities interact with participation logic, governance maturity, and long-term stewardship in operational settings.

3.2.7 Distributed Knowledge Nodes and Domain-Level Exchange

Taken together, the RFI contributions indicate a structural possibility that extends beyond a single Knowledge Manager deployment. The material suggests an innovation commons system composed of distributed knowledge nodes that exchange domain-specific information through shared institutional infrastructure.

In such a configuration, knowledge remains embedded within organisations and communities operating in different institutional contexts, sectors, and geographic locations. Each node develops expertise shaped by its local environment, regulatory conditions, and practical experience. Even when actors operate within the same knowledge domain, their perspectives and practices therefore tend to diverge over time.

This diversity is not necessarily a weakness. It may instead constitute a primary source of value in an innovation commons system.

When domain-specific knowledge nodes exchange information under shared governance arrangements, the ecosystem gains access to a broader pool of distributed information. Practices developed in one setting become visible to others. Methods, tools, and interpretations that emerged under one set of conditions can be examined and adapted in another.

Technically, such exchange requires infrastructure capable of linking artefacts produced in different organisational environments. The RFI responses described several architectural mechanisms that make this possible.

Redpill Linpro highlighted semantic interoperability as a prerequisite for distributed knowledge reuse. Consistent metadata models, machine-readable tagging, and open indexing allow artefacts produced in one node to become discoverable and interpretable in another.

AI-Bees approached the problem through federation. Their architecture introduces AI pods as autonomous execution and data domains in which each organisation retains local control over storage, models, and access policies. Knowledge exchange occurs through policy-driven interfaces supported by verifiable identities and certificate-based roles.

Sweco focused on the architectural mechanisms required for interoperability. Modular services, API-based integration, and contract-based interfaces allow organisational nodes to evolve independently while remaining connected through shared technical standards and traceable infrastructure pipelines.

Taken together, these elements describe a system in which domain knowledge can circulate across organisational nodes while remaining institutionally anchored where it was produced.

This structure becomes particularly relevant when similar knowledge domains exist across multiple organisations or regions. Actors working within the same domain—such as digital health, municipal services, energy transition, or public-sector digitalisation—often accumulate experience through local experimentation. Each node develops partial knowledge shaped by its own regulatory environment, technical infrastructure, and operational priorities.

When these nodes are connected through interoperable institutional infrastructure, their accumulated experience becomes collectively visible. Practices developed in one node can inform experimentation in another. Artefacts produced under one set of conditions may be recombined with approaches developed elsewhere. Over time, this exchange allows domain knowledge to evolve through recombination rather than through central coordination.

The potential value of such exchange increases when knowledge nodes connect beyond local ecosystems. National networks may link municipalities, agencies, and research environments working within the same domain. At a larger scale, similar infrastructures may enable international exchange between actors facing comparable challenges in different institutional settings. In both cases, the system enables domain knowledge to circulate without requiring central ownership or uniform organisational structures.

Several forms of value creation may emerge from such distributed exchange:

- emerging practices within a domain become visible across nodes
- solutions developed under different institutional conditions can be compared
- artefacts and methods can be recombined across organisational boundaries
- shared conceptual frameworks and metadata structures may gradually stabilise

Technically, this type of ecosystem interaction depends on alignment across three layers:

- semantic structures that allow artefacts to be interpreted across nodes
- governance arrangements defining licensing conditions and role-based responsibility
- modular technical architecture enabling interoperability without tight system coupling

The Knowledge Manager explored in WP3 can be interpreted within this structure as an orientation mechanism. By linking artefacts, metadata, licensing conditions, and traceable sources, it helps actors navigate a distributed knowledge landscape and identify relevant contributions produced by other nodes.

If such infrastructures connect nodes across organisations, regions, or national contexts, the innovation ecosystem gains the ability to circulate domain knowledge while preserving institutional autonomy. The innovation commons resource then emerges not from central accumulation but from the structured interaction between distributed knowledge nodes operating within shared governance arrangements.

3.2.8 Concluding Reflection: What the Dialogue Added

The joint discussion with respondents and the project group added important clarifications to the written RFI material.

The MVP was explicitly described as a constrained proof of concept. It implemented a chat-based interface over indexed documents, with emphasis on file handling, traceability, licensing conditions, and access control rather than advanced AI capability. The Azure deployment reflected speed of experimentation rather than long-term architectural commitment. This reinforces the distinction between technical feasibility and institutional readiness.

The discussion also made licensing conditions more operational. Documents in the prototype carry explicit licence metadata that is surfaced in generated answers. This indicates an approach where licensing conditions are embedded in the institutional infrastructure rather than treated as external legal artefacts. Traceability in outputs depends directly on traceability in source artefacts and their version history.

Participants further clarified that the Knowledge Manager should be understood as a component within a broader innovation commons community. The prototype intentionally excluded member-to-member interaction. This makes visible a structural boundary: the Knowledge Manager supports distributed information retrieval but does not itself constitute the collective action layer required for a functioning innovation commons community.

The architectural discussion reinforced the importance of modular pipelines, API-based integration, and platform independence. MCP was highlighted as a potential standard for connecting models and tools, supporting replaceability over time. The dialogue also suggested that model size is often less decisive than pipeline design and data quality.

Digital sovereignty emerged as a live institutional concern rather than an abstract preference. Interest in European cloud alternatives was framed in terms of resilience and long-term

control. At the same time, participants noted that specialised and smaller models may often be sufficient, provided that traceability and governance arrangements are robust.

Finally, the discussion introduced a methodological caution regarding early formalisation. Graph models and ontologies were recognised as powerful but potentially constraining in domains characterised by entrepreneurial discovery and evolving concepts.

Overall, the dialogue indicates that the long-term viability of an innovation commons system depends not only on modular architecture and distributed trust, but on sustained stewardship capable of maintaining traceability, licensing conditions, and adaptive governance arrangements over time.

3.2.9 Synthesis

Taken together, the three RFI contributions outline a possible architecture for an innovation commons system.

Redpill Linpro clarified participation logic through relevance, quality assurance, and semantic accessibility. AI-Bees addressed distributed trust and federated governance arrangements. Sweco operationalised these principles through modular, contract-based, and sovereignty-oriented infrastructure.

The material suggests an emerging alignment between participation logic, governance arrangements, and technical architecture. Several questions remain indicative rather than conclusive:

- how contribution patterns evolve over time
- what organisational capacity trust frameworks require
- how marketplace incentives affect sustainability
- what level of sovereignty is sufficient in practice
- how stewardship remains coherent in distributed systems

3.3 Building and Testing the Knowledge Manager

If the RFI clarified the necessary design layers for a Knowledge Manager in an innovation commons, the development of the MVP allowed us to examine what happens when these principles encounter operational reality.

The objective of this phase was not to build a finished product. It was to test whether the combined logic of incentives, governance, and architecture could be translated into a functioning knowledge service embedded in a real community.

The MVP therefore functioned as a probe rather than a solution. It enabled us to explore how governance principles become encoded in system design, how architectural decisions shape institutional behavior, and where technical feasibility diverges from collective value creation.

The work in WP3 built on earlier insights from the pre-study Information Structures for Learning Networks (ISLNd), which demonstrated that AI-supported knowledge management requires structured information models, clearly defined responsibilities, and governance alignment in order to generate usable value. WP3 extended this insight by moving from conceptual exploration to applied piloting within an established innovation commons.

Rather than presenting the development process chronologically, this section focuses on what became visible through implementation: the intended role and scope of the Knowledge Manager, the operationalization of governance in technical architecture, and the structural constraints revealed through testing.

3.3.1 Purpose and scope of the Knowledge Manager

The Knowledge Manager was developed with a specific infrastructural hypothesis: that an AI-based knowledge service can strengthen the collective capacity of an innovation commons if — and only if — it is aligned with participation logic, governance principles, and architectural robustness.

Its intended role was not to replace community dialogue, nor to function as an autonomous decision-making system. Instead, it was designed to support three interconnected capacities within the commons:

- Lowering the threshold for accessing accumulated knowledge
- Supporting onboarding and orientation for new participants
- Enabling systematic reuse of existing artifacts
- Strengthening the long-term collective memory of the network

These objectives directly reflect the design layers identified in the RFI:

- From an incentive perspective, the KM should make contribution visibly valuable through reuse and accessibility.
- From a governance perspective, it should respect roles, responsibilities, and licensing conditions embedded in the commons.
- From an architectural perspective, it should be modular, traceable, and adaptable.

Demand Acceleration (DA) and its associated community were used as a case context for the MVP. The system was trained and tested exclusively on DA-related artifacts, primarily structured and human-curated documents such as case reports, guides, and methodological descriptions. This deliberate scoping decision reflected a strategic choice: to prioritize reliability, traceability, and governance alignment over breadth of data sources.

By narrowing the scope to curated artifacts, the project could focus on testing institutional alignment rather than maximizing data ingestion. This choice later proved significant when evaluating system behavior and quality.

3.3.2 Technical approach and architecture

From a technical perspective, the Knowledge Manager was built on a retrieval-augmented generation (RAG) architecture. However, the architectural choices were not driven primarily by model performance. They were guided by the need to operationalize governance, traceability, and institutional alignment.

The system processes documents through a structured pipeline consisting of ingestion, chunking, vectorization, indexing, and retrieval. A chat-based interface enables natural-language interaction with the indexed corpus. While this setup reflects prevailing best practices for AI-supported knowledge services, its design was deliberately shaped by the three infrastructural layers identified earlier: incentives, governance, and architecture.

Architecture as Institutional Encoding

Several architectural decisions were made explicitly to embed governance principles into the system:

1. Document Container Logic and Versioning

Each artifact was treated as a logical container to which one or more document versions could be attached. This ensured traceability over time and clarified responsibility for updates. The container logic reinforced the idea that artifacts in a commons are not static files but evolving institutional objects.

2. Separation of Content, Metadata, and Access Logic

The architecture maintained a clear separation between:

- The content itself (documents and files),
- Metadata describing ownership, licensing, and status,
- Access and visibility rules governing who can view or manage artifacts.

This separation was essential to ensure that governance rules could evolve without requiring structural redesign of the system.

3. Role-Based Administration and Limited Visibility

Administrative functionality was intentionally restricted. Uploading, publishing, unpublishing, and licensing management were not open actions but governed by defined roles. This design choice reflected the insight that quality assurance and stewardship cannot be fully automated.

4. Licensing Integration

Licensing was not treated as external documentation but as structured metadata directly associated with each artifact. The system was designed to inform users about applicable conditions when referencing documents, reinforcing transparency in reuse.

Exploration of Alternative Technical Paths

Before consolidating the MVP architecture, several alternative approaches were explored:

- Evaluation of existing AI-supported knowledge platforms combined with external document repositories.
- Hands-on prototyping (“vibe coding”) to deepen internal understanding of embedding strategies, chunking logic, and vector database behavior.
- Collaboration with an external consultant to develop an Azure-based extended MVP with repository management and administrative control.

These explorations confirmed that technical feasibility is not the primary bottleneck. Mature cloud infrastructures, vector databases, and RAG pipelines are sufficiently robust for applied implementation.

The decisive challenge lies elsewhere: ensuring that architectural design remains aligned with governance logic and participation dynamics.

Architectural Trade-offs and Constraints

Several trade-offs became apparent during implementation:

- Greater flexibility in data ingestion increases complexity in governance control.
- More granular chunking improves retrieval precision but complicates traceability and attribution.
- Full automation reduces administrative burden but weakens stewardship mechanisms.

These tensions reinforced a central methodological insight: in innovation commons, technical architecture is not neutral infrastructure. It actively shapes how knowledge is structured, attributed, and reused. For this reason, the MVP prioritized modularity, traceability, and governance alignment over maximum automation or feature breadth.

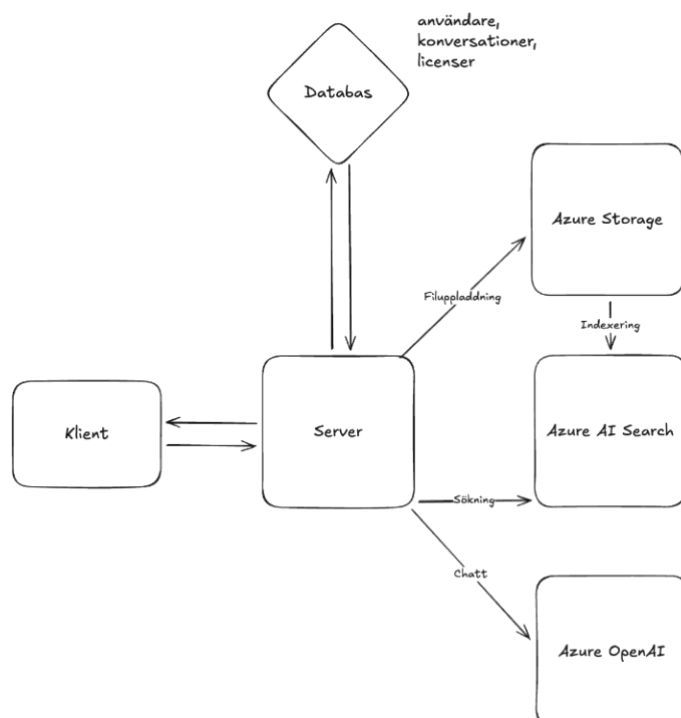


Figure 1 Knowledge Manager MVP architecture

3.3.3 Testing, evaluation, and technical limitations

Evaluation Design

The Knowledge Manager was evaluated as an institutional prototype rather than a finished system. The MVP functioned as a hypothesis-testing probe within a complex adaptive socio-technical system rather than as a delivery instrument. The objective was to understand how AI-supported retrieval and synthesis function within a governance-sensitive innovation commons environment.

To strengthen analytical validity despite a limited number of structured user tests, two complementary evaluation mechanisms were applied:

- **Structured user testing**, based on realistic role-based scenarios (Process Leader and Procurement Lead).
- **Behavioural log analysis**, examining actual interaction patterns across sessions.

This combination enabled triangulation between *reported experience* and *observed use behaviour*.

Finding 1 – Structured Orientation and Drafting Support

Observations from user testing

Users consistently described the Knowledge Manager as particularly helpful in structured, method-oriented tasks. Outputs were experienced as clear, concrete and practically reusable with limited editing effort. Drafting support — such as structured explanations of Demand Acceleration or preparation of presentations — was highlighted as valuable.

Observations from log analysis

The interaction logs reinforce this pattern. A substantial proportion of questions revolve around:

- Demand Acceleration definitions
- RFI design
- Procurement principles
- Evaluation criteria

Conversations are typically short and task-oriented, indicating point-of-need use rather than exploratory dialogue.

Integrated analysis

Both data sources suggest that the Knowledge Manager currently functions most robustly as a *cognitive compression and structuring tool*. Its primary value lies in transforming institutional complexity into communicable structure.

It is less used as a discovery mechanism and not yet as a sustained collaborative reasoning partner.

Finding 2 – Credibility Depends on Institutional Precision

Observations from user testing

Trust was closely linked to specificity and correctness. Where responses were precise and aligned with procurement logic, users expressed confidence. Minor inaccuracies (e.g., incomplete articulation of procurement principles) were perceived as potentially undermining credibility.

Observations from log analysis

The logs show repeated foundational questions (e.g., “What is Demand Acceleration?”), sometimes across separate sessions. This indicates either continued conceptual ambiguity among users or insufficient consolidation of core definitions within the system.

Meta-questions regarding accessible documents also appear, suggesting users actively test epistemic boundaries.

Integrated analysis

In governance-sensitive contexts, rhetorical fluency is insufficient. Institutional precision and conceptual consistency are decisive for legitimacy.

The findings suggest that definitional stability and metadata discipline are not technical details, but core governance functions of the Knowledge Manager.

Finding 3 – Governance Architecture Is Institutionally Central but Interactionally Invisible

Intended focus

A deliberate design ambition in WP3 was to surface governance mechanisms — including licensing and usage conditions — as part of the Knowledge Manager architecture. The system was structured to link artefacts, roles, and usage conditions in order to operationalise commons principles.

Observations from user testing

No user explicitly commented on licensing or usage conditions.

Observations from log analysis

Licensing logic and usage conditions do not appear as explicit themes in user interactions. Questions focus overwhelmingly on operational, procedural, and conceptual matters.

Integrated analysis

The absence of feedback on licensing is analytically significant.

Possible interpretations include:

1. Governance functions were perceived as background infrastructure rather than user-facing value.
2. Users prioritised operational task support over rule frameworks in early interaction.
3. Licensing logic is not yet cognitively integrated into everyday professional reasoning.

This does not imply that licensing is irrelevant. On the contrary, licensing remains foundational from a commons governance perspective. However, the testing suggests that governance architecture does not automatically become interactionally salient.

This creates a structural design tension:

- Increasing visibility may increase cognitive load.
- Keeping governance infrastructural risks invisibility of institutional robustness.

The evaluation indicates that licensing must be assessed not only for legal coherence but for interactional relevance and timing.

Finding 4 – Technical Stability and Retrieval Precision Affect Trust

Observations from user testing

Occasional system interruptions (“cannot generate answer”) were reported. Even if limited, such instability may disproportionately affect trust in institutional contexts.

Observations from log analysis

The logs include instances of:

- Retrieval limitations
- Inability to provide information despite documentation being available
- Predominantly short interaction cycles

These patterns indicate that the system is sensitive to metadata structure and retrieval logic.

Integrated analysis

Technical reliability is inseparable from institutional legitimacy. Retrieval precision, stable runtime behaviour, and clear scope boundaries are prerequisites for sustained reuse in a commons environment.

Short interaction cycles further suggest that the Knowledge Manager is currently positioned as a transactional support tool rather than an embedded collaborative infrastructure.

Overall Maturity Assessment

Across findings, a coherent maturity profile emerges.

Confirmed strengths

- Strong value in structured orientation and drafting.
- Clear usefulness in DA- and RFI-related professional contexts.
- Curated artefacts provide a stable knowledge base.

Partially validated

- Role-based relevance is evident, but UX friction and instability reduce depth of engagement.
- Conceptual consistency requires active curation.

Structurally unresolved

- Governance mechanisms are architecturally central but not interactionally visible.
- Sustained iterative co-creation has not yet emerged as a usage pattern.
- Metadata discipline remains decisive for scaling reliability.

Concluding Insight

Testing clarifies not whether the AI “works,” but what institutional role it can credibly occupy at this stage.

The Knowledge Manager currently functions most robustly as:

- An onboarding accelerator
- A structured drafting assistant
- A compliance-adjacent decision support tool

It does not yet function as:

- A collaborative knowledge co-creation environment
- A substitute for governance
- A self-sustaining commons mechanism

The central conclusion is institutional:

Technical feasibility is demonstrable.

Institutional maturity — including precision, metadata discipline, governance integration, and stewardship — determines long-term viability.

3.3.4 Licensing, conditions, and knowledge commons

One of the most significant insights from WP3 concerns **licensing and conditions for shared knowledge artifacts**. As the Knowledge Manager produces responses that typically draw on multiple documents, it becomes difficult—if not impossible—to attribute individual parts of an AI-generated answer to specific source documents and their respective licenses.

This challenge is not only one of attribution in a practical sense, **nor is it unique to innovation commons**. Similar tensions arise whenever information is extracted, synthesised, restructured, or circulated across multiple artefacts and contexts. Current debates on the training of large language models provide one visible example, but related issues also appear in AI-generated summaries, search indexing, embeddings, knowledge graphs, and more conventional forms of synthesis such as briefs, presentations, and composite guidance documents. In WP3, the issue becomes visible because the **infrastructure manages artefacts**, while AI-supported synthesis increasingly operates on informational content that **cuts across multiple artefacts**. The resulting tension points to a distinction between artefact-level management and information-object-level governance, the latter being developed more fully in the conceptual framework of the final report. It also suggests that AI-supported Knowledge Managers are only one possible mechanism for sharing and reusing information, knowledge, and insights within an innovation commons system, rather than the only or necessarily dominant one.

To address this, WP3 expanded the concept of licensing beyond traditional Creative Commons frameworks to include three distinct dimensions:

- Copyright and licensing (e.g. Creative Commons),
- Rights related to use and reuse of the content (IP),
- Rights related to dissemination beyond the community.

Based on this, a set of **DA-specific conditions** was defined and implemented in the KM. Each document in the repository was associated with a specific condition, and the KM was designed to inform users about applicable conditions when referencing documents.

At the same time, this work highlighted a fundamental challenge: **AI-based synthesis blurs the boundaries between individual artifacts**. One implication is that innovation commons using AI-based knowledge services may need to adopt harmonized conditions across repositories to avoid legal and ethical ambiguity.

3.3.5 Relation to the community platform

WP3 also included initial work on integrating the Knowledge Manager with a community platform (Mighty Networks). While the technical and social integration between community, platform, and KM was not fully realized during the project period, the work clarified the importance of viewing these components as a **socio-technical system** rather than standalone tools.

The Knowledge Manager should be understood as a complement to—rather than a replacement for—human interaction, dialogue, and collective sense-making within the community.

3.3.6 Summary of key learnings

WP3 explored the Knowledge Manager as an institutional prototype within an innovation commons system. The work did not assess technical performance in isolation, but examined how participation logic, governance arrangements, and architectural design interact under operational conditions.

The findings suggest four main conclusions.

1. AI-supported knowledge services can create practical value when grounded in curated artefacts.

Within the tested scope, the Knowledge Manager functioned most clearly as support for orientation, onboarding, and structured drafting. This suggests that AI-based retrieval and synthesis can reduce cognitive friction and improve access to distributed information when metadata, traceability, and definitional precision are sufficiently robust.

2. Institutional precision is a precondition for legitimacy.

Technical feasibility alone was not sufficient to generate trust. User confidence was closely linked to conceptual consistency, retrieval precision, and clarity regarding scope and limitations. In governance-sensitive contexts, even minor inaccuracies risk undermining credibility.

3. Governance mechanisms proved structurally central but interactionally weakly visible.

Licensing conditions and role-based responsibility were embedded in the architecture, yet they rarely became explicit themes in use. This suggests that governance arrangements may initially function as background infrastructure rather than user-facing value, even when they remain decisive for long-term viability.

4. Technical deployment did not in itself generate collective action.

Observed use was mainly short, task-oriented, and transactional. The Knowledge Manager did not yet function as a collaborative co-creation environment or as a substitute for stewardship within the innovation commons community. This indicates that AI can support pooling, but cannot replace the institutional conditions required for sustained participation.

Taken together, these findings indicate that the main constraint in developing AI-supported knowledge services for innovation commons is not model capability as such, but institutional maturity. Metadata discipline, stewardship, governance arrangements, and traceability shape the conditions under which AI can strengthen rather than distort the shared resource.

At the same time, one important area remains only partially explored. The current work prioritised formal, human-curated artefacts, while a substantial share of knowledge creation takes place in meetings, dialogue, and ongoing interaction. **Conversational material may therefore represent an important extension** of the innovation commons resource, but its integration would require further clarification of licensing conditions, stewardship responsibilities, and traceability mechanisms.

WP3 should therefore be understood as an institutional learning process rather than a finished solution. The results suggest that AI-based Knowledge Managers can contribute to innovation commons systems under specific conditions, but that long-term viability depends on continued alignment between architecture, governance arrangements, participation logic, and community stewardship.

4. PART II – Exploring the Community Platform

4.1 Why Technology Alone Was Never Enough

In the beginning of our test journey, much attention was directed toward the Knowledge Manager and the technical possibilities of AI-supported knowledge services. However, it quickly became clear that even the most well-designed technical system would be insufficient if it was not embedded in a functioning community.

The underlying assumption of the project was never that technology could replace community. Rather, the hypothesis was that technology might support and strengthen an innovation commons. This distinction proved critical.

An innovation commons is not primarily a repository of artefacts. It is a living socio-technical system. It depends on interpretation, negotiation, trust, shared language, and evolving norms. These dimensions cannot be automated.

AI can assist with:

- structuring information
- lowering access thresholds
- retrieving artefacts
- supporting onboarding

But AI cannot:

- establish legitimacy
- resolve conflicts
- define shared purpose
- sustain voluntary participation
- carry responsibility

Ostrom's design principles for robust commons emphasise clearly defined boundaries, participatory rule-making, monitoring, conflict-resolution mechanisms, and nested governance structures. None of these are reducible to software features. They require active stewardship.

The community platform was therefore not conceived as a communication tool alone. It was understood as an institutional space where roles, expectations, and norms could be enacted in practice. Without this layer, the Knowledge Manager risked becoming an isolated technical artefact detached from the social processes that give knowledge meaning.

This insight reframed WP3. The challenge was not to build a better tool. It was to understand how digital infrastructure interacts with human coordination.

Technology was necessary. It was never sufficient.

4.2 What Happened When People Actually Used It

The community platform (Mighty Networks) was introduced gradually within the Demand Acceleration network. Rather than launching a large-scale public rollout, the platform was integrated into existing relational structures.

This approach allowed us to observe behaviour patterns under relatively controlled conditions.

Several recurring dynamics became visible.

1. Engagement followed a power-law distribution

A small group of highly active participants generated a disproportionate share of activity. Most members remained passive observers. This pattern is well-documented in community research and was confirmed in practice.

Average engagement metrics therefore provided limited insight. The vitality of the platform depended heavily on a small number of stewards and committed contributors.

2. Facilitation mattered more than features

Technical functionality did not automatically produce activity. Structured prompts, curated discussions, and deliberate facilitation were necessary to sustain momentum.

When facilitation decreased, activity declined.

This suggests that platforms do not create community; they amplify or dampen existing relational energy.

3. Administrative efficiency improved significantly

Before the platform, community coordination relied on email lists, manual updates, and fragmented communication channels. The platform reduced friction in:

- onboarding participants
- distributing information
- managing events
- maintaining contact lists

This operational efficiency is easy to underestimate but proved valuable in practice. Even when engagement levels fluctuated, the platform provided structural stability.

4. Adoption friction remained real

Despite clear value, behavioural change was limited. Many participants did not integrate the platform into daily routines. Notification overload and competing digital environments reduced attention.

This highlights an important limitation: digital commons must compete within saturated attention economies.

5. Some failures were quiet

Certain discussion threads did not gain traction. Some features were rarely used. The absence of explicit resistance should not be interpreted as success.

Silence in a community platform is analytically ambiguous. It may indicate satisfaction, indifference, overload, or lack of perceived relevance.

6. Human curation increased quality

Content that was curated, framed, or summarised generated more meaningful engagement than unstructured postings. This reinforced insights from the Knowledge Manager work: structure increases usability.

4.3 What We Learned About Communities and Commons

The experience with the community platform clarified several distinctions that are central to innovation commons governance.

Sharing is not the same as co-creating

Uploading information is technically easy. Co-creating meaning requires interaction, feedback, and iterative refinement. The latter demands time, trust, and perceived reciprocity.

A platform can enable sharing. It does not automatically enable co-creation.

Participation differs from stewardship

A commons depends on active stewards. Without individuals who curate, moderate, and nurture norms, shared resources degrade or stagnate.

Digital infrastructure reduces coordination costs but does not eliminate the need for stewardship. If anything, it makes stewardship more visible and more critical.

Infrastructure shapes behaviour

Platform design influences interaction patterns. Features such as visibility of posts, reaction mechanisms, moderation rights, and group segmentation affect how knowledge flows.

Technical architecture therefore becomes a governance choice. The way permissions, visibility, and access are structured encodes institutional logic.

Openness must be balanced with clarity

An innovation commons aspires to openness. Yet excessive openness without defined expectations can reduce contribution quality and participant confidence.

Clear boundaries, role definitions, and shared norms increase participation rather than restrict it. This mirrors Ostrom's findings in traditional commons contexts.

Neglecting community undermines technical investments

The Knowledge Manager depends on structured artefacts. Structured artefacts depend on contributors. Contributors depend on community legitimacy and perceived value.

This creates a circular dependency:

Community → Artefacts → AI Value → Community Reinforcement

Breaking any link weakens the system.

Interim Conclusion

The exploration of the community platform confirms that innovation commons are fundamentally socio-technical systems.

Digital infrastructure can:

- reduce coordination costs
- increase transparency
- support collective memory
- support monitoring and evaluation of participation and use

But sustainable knowledge commons require:

- stewardship
- facilitation
- governance clarity
- human judgment

The community platform did not solve these challenges. It made them visible.

In doing so, it provided essential insight for the continued development of hybrid infrastructures where AI and community are designed together rather than treated as separate domains.

4.4 Research context: participation and value creation in community platform

Research on commons and digital community platforms points to a relatively consistent picture: sustained participation is shaped primarily by perceived value, trust, and effective social organization, rather than by technical functionality alone. Platforms may enable interaction, but stable activity emerges only when participants perceive the environment as relevant, legitimate, and worth their time investment (Ostrom, 1990; Hess & Ostrom, 2007).

Jason Potts further develops this perspective through the concept of the innovation commons, in which communities are understood as institutional arrangements for pooling dispersed knowledge under conditions of uncertainty. In early stages of innovation, value is often unclear, making commons-like collaboration particularly important for enabling collective exploration and the discovery of new opportunities (Potts, 2018; Potts, 2019). From this perspective, the community platform becomes not merely a meeting place but a structure that supports collective value creation.

Empirical studies from the 2020s confirm that participation in online communities is typically driven by a combination of intrinsic motivations (such as meaningfulness and perceived competence), social factors (recognition and belonging), and more instrumental benefits related to work or problem solving. Recognition from other users can strengthen willingness to contribute, while negative social signals may have the opposite effect (Mustafa et al., 2022). Several studies also show that identification with the community's purpose is a strong predictor of active participation.

At the same time, the utility perspective emerges as central. Community platforms tend to retain participants when they deliver practical usefulness, such as faster problem solving, access to relevant experiences, or support for professional development (Al-Khasawneh et al., 2023; Hwang & Foote, 2021). Platform value thus arises through the interplay between social dynamics and practical usefulness.

A recurring theme in the literature is that community platforms manage heterogeneous artefacts of varying maturity and quality. Contributions may range from brief questions and comments to more developed guides, templates, or structured resources. Studies of professional and educational communities show that participants' motivation and behavior are closely related to the quality of the artefacts produced, underscoring the need for mechanisms for feedback, visibility, and curation (Gutiérrez-Páez et al., 2023). Platforms that function well over time therefore tend both to lower the threshold for small contributions and to create conditions that allow more advanced contributions to be developed and reused.

Research also indicates that knowledge domain and context influence participation dynamics. In professional networks, participation is often closely tied to practical work relevance, whereas more mission-driven or volunteer communities are more strongly motivated by identification with the cause and prosocial engagement. In innovation and co-creation environments, the opportunity to influence development emerges as a central driver.

Overall, the literature suggests the need for domain-sensitive approaches rather than one-size-fits-all engagement models.

Against this background, a clear practical implication emerges: active facilitation over time is a key success factor. Studies show that ongoing community leadership—through encouragement, visibility of contributions, curation, and relationship building—can have measurable effects on participation and activity (Hornuf & Jeworrek, 2023; Sharp, 2024). The role functions as a link between platform structure and participant motivation and contributes to maintaining direction and quality in the community.

This aligns closely with our own experiences from working with our community platform. Our assessment is that technical functionality constitutes a necessary foundation, but that continuous community management is critical for sustaining engagement, direction, and quality over time. In the following section, we elaborate on these experiences and how they can be translated into a more systematic approach.

4.5 The Platform as a Monitoring and Evaluation Infrastructure

The community platform does more than host interaction. It generates a continuous stream of behavioural data that, if attended to systematically, can serve as a central component in a monitoring and evaluation framework for an innovation commons system.

Conventional approaches to evaluation tend to treat it as a discrete activity: an assessment conducted at a fixed point to determine whether targets have been met. In a living innovation commons community, this logic breaks down. Who participates, how frequently, in what ways, and around what information assets -- these patterns shift continuously, and the signals they produce are most useful precisely when they can be read in real time and fed back into decisions about facilitation, governance, and the further development of the resource pool. What is needed is not periodic reporting but embedded, ongoing attention to what is actually happening.

A digital platform is structurally well suited to this. At the most basic level, platform data tracks membership growth, activity patterns, contributions of information assets, and engagement with specific parts of the resource pool over time. A declining engagement curve is not simply an operational concern; it is a signal that assumptions about perceived value or relevance may need revisiting. A concentration of activity around a narrow set of contributors raises questions about stewardship and the long-term sustainability of the innovation commons system. The platform makes these dynamics visible in ways that informal observation cannot.

Beyond passive observation, a digital platform enables active inquiry. Surveys can be distributed to specific segments of the community, administered at particular moments in a community lifecycle, or repeated over time to track shifting perceptions. User tests can be embedded directly in the platform -- observing how community members navigate information assets, where they encounter friction, and what they consistently ignore. Silence in a community platform is analytically ambiguous, but with the right instrumentation it becomes interpretable rather than invisible.

Our own experience with the community platform has so far only scratched the surface of what this infrastructure makes possible. Usage patterns were tracked and informally interpreted, but a coherent evaluation framework was never fully operationalised. What the experience suggests, however, is the potential: with appropriate instrumentation and analytical discipline, platform data could support a learning-oriented approach to evaluation that runs continuously through the lifecycle of an innovation commons system rather than being applied episodically or after the fact.

This points to a design implication that extends beyond the current project. A community platform built to support an innovation commons system should, from the outset, be designed with evaluation in mind. That means attending to what data is generated and how it is structured; what questions can be explored through surveys or structured interaction; and how insights from platform behaviour can be fed back into collective decision-making about governance and the development of shared information assets. The platform is not merely a venue. Handled with intention, it is an observatory.

5. PART III – When AI Meets Community

5.1 The Interplay: Where the Real Value Emerges

Part I and Part II examined two interdependent components of an innovation commons system: the Knowledge Manager and the community platform. The updated analysis in Part I adds a third structural element: distributed knowledge nodes and domain-level exchange. Together, these three components clarify that value emerges not from a single tool, but from how artefacts, communities, and governance arrangements interact across levels.

The Knowledge Manager reduces friction in access, orientation, and structured reuse of curated artefacts. The community platform sustains participation logic through facilitation, stewardship, and collective sense-making. Distributed knowledge nodes introduce an additional capability: domain-level resource formation and exchange without full centralisation. This configuration suggests a layered system in which pooling occurs locally where context is strongest, while exchange occurs across domains where interoperability, traceability, and licensing conditions remain coherent.

The interplay therefore concerns two complementary alignments:

- **Vertical alignment (within a node):** community practices → curated artefacts → AI-supported retrieval
- **Horizontal alignment (across nodes):** shared governance arrangements → traceability → domain-level exchange

Two failure modes become visible when one layer is missing.

AI without community and stewardship may produce technically fluent outputs with weak legitimacy and limited long-term viability.

Community without structured artefacts and traceability may generate rich interaction but little sustained reuse, making pooling fragile and difficult to scale.

A third failure mode is now also relevant:

Nodes without interoperable exchange may stabilise local value but fragment the overall innovation commons system into parallel micro-commons with limited learning transfer.

The findings suggest emerging design principles for hybrid infrastructures:

- AI should support sense-making and reuse, not substitute for stewardship.
- Domain-level exchange requires disciplined interoperability: metadata, traceability, and licensing conditions.

- Governance must be encoded in architecture, but remains dependent on human judgement and institutional conditions.
- Long-term viability depends on alignment across levels: within nodes and between nodes.

5.2 What This Means for Governance and Sustainability

The report indicates that technical choices function as governance choices at two levels: within communities and across distributed knowledge nodes. Infrastructure shapes what becomes visible, reusable, and legitimate. It also shapes boundaries: who participates, what can be pooled, and under what licensing conditions.

Introducing distributed nodes strengthens sustainability in one dimension: it can preserve context and organisational autonomy while enabling contribution. At the same time, it increases governance complexity. Domain-level exchange requires shared rules that remain interpretable across different institutional contexts. This makes traceability and licensing conditions not just repository features, but system-level governance mechanisms.

From a sustainability perspective, three implications follow:

1. **Stewardship is the binding constraint.** The more distributed the system, the more critical the role of stewards who can maintain metadata discipline, quality assurance, and governance consistency over time.
2. **Interoperability becomes an institutional condition.** Domain-level exchange depends on shared semantics, compatible artefact structures, and predictable licensing conditions. Without this, the system risks becoming a set of disconnected nodes.
3. **Evaluation must follow the system structure.** What counts as value differs by level: local usability within a node, and learning transfer/reuse across nodes. This implies that evaluation should monitor both node-level performance and cross-node exchange quality, rather than relying on single platform metrics.

These implications also inform how WP2 may interpret digital infrastructure. The community platform, the Knowledge Manager, and distributed knowledge nodes should be treated as parts of institutional infrastructure, not as standalone tools.

5.3 What We Would Recommend (and What We Would Not)

The work does not support a universal blueprint. Different starting points and institutional conditions create different viable paths. The findings do, however, support grounded principles.

Recommendations

1. **Design for a multi-level innovation commons system.** Treat nodes as meaningful units of governance and context. Use domain-level exchange to enable learning transfer without forcing centralisation.
2. **Prioritise traceability and licensing conditions as system-level functions.** If knowledge is exchanged across nodes, governance arrangements must remain coherent across contexts. This requires disciplined metadata and explicit responsibility structures.
3. **Maintain a clear division of labour between AI and community.** Use AI for retrieval, synthesis, onboarding, and drafting support; rely on community practices for legitimacy, stewardship, and collective action.
4. **Build stewardship capacity as core infrastructure.** Stewardship cannot be treated as residual work. It is the mechanism that stabilises quality, participation logic, and long-term viability across nodes.
5. **Use prototypes as hypothesis-testing probes.** Continue to treat implementations as learning instruments, especially for cross-node governance and exchange.

What we would not recommend

- Centralising all artefacts without a clear governance logic for participation and reuse
- Treating interoperability as a technical integration task rather than an institutional condition
- Assuming that exchange across nodes will work without shared semantics and licensing discipline
- Assuming that AI capability can compensate for weak stewardship or unclear governance arrangements

6. Closing Reflections: From Pilot to Possibility

This report set out to examine whether **digital infrastructure** can support **innovation commons systems**, and under what **institutional conditions**. Across WP3, the answer appears neither simply affirmative nor simply negative. The findings suggest that digital infrastructure can make important contributions, but only when it is treated as part of an **institutional infrastructure** rather than as a standalone technical solution.

Part I showed that the **Knowledge Manager** can create practical value where the innovation commons resource consists of curated artefacts with sufficient **traceability**, metadata discipline, and conceptual stability. **Part II** showed that **technology alone does not sustain participation**: community practices, facilitation, and **stewardship** remain decisive for collective action and **long-term viability**. **Part III** extended this perspective by indicating that **distributed knowledge nodes** may support domain-level resource formation and exchange, but only where governance arrangements, interoperability, and licensing conditions remain coherent across contexts.

Taken together, the report points to a consistent conclusion. The central challenge is not primarily how to build more advanced tools, but how to align **architecture, governance arrangements, participation logic**, and **stewardship** over time. AI-supported retrieval may reduce friction. Community platforms may support participation. Distributed nodes may preserve context while enabling exchange. None of these components, however, removes the need for institutional conditions that make **pooling** sustainable.

The work also makes visible a further distinction that becomes increasingly important as **AI-supported synthesis** develops. The infrastructure examined here operates primarily at the level of **artefacts and repositories**, because that is the level at which documents, metadata, traceability, and licensing conditions can be managed in practice. Yet some governance challenges arise at the level of **informational content** that moves across multiple artefacts. This is not unique to innovation commons, nor to Knowledge Managers. It appears more broadly wherever information is extracted, synthesised, restructured, and reused across contexts. In WP3, this tension became visible most clearly in relation to licensing and attribution.

Several questions therefore remain unresolved. The report does not establish how **conversational material** can be integrated responsibly at scale. It does not resolve how stewardship should be organised and financed over longer periods. It does not yet show how **cross-node exchange** behaves under substantially larger, more diverse, or more conflict-prone conditions. Nor does it remove the **structural tension** between project-based development logic and the long-term maintenance needs of innovation commons systems.

For these reasons, the work should be understood as a **starting point rather than a conclusion**. It narrows the field of plausible assumptions, clarifies structural tensions, and identifies conditions under which digital infrastructures may contribute to reuse, collective action, and long-term viability. The next step is therefore not simply further implementation, but continued experimentation with explicit attention to governance, stewardship, traceability, interoperability, and the relation between **artefact-level infrastructure** and **information-object-level governance**.

References

- Al-Bees. (2025). *Federated innovation commons: Response to request for information (RFI)*. Development and Application of Innovation Commons for Strengthening Innovation Eco-Systems project document.
- Al-Khasawneh, M., Al-Haddad, S., Sharabati, A.-A. A., Al Khalili, H. H., Azar, L. L., Ghabayen, F. W., Jaber, L. M., Ali, M. H., & Masa'deh, R. (2023). How online communities affect online community engagement and word-of-mouth intention. *Sustainability*, 15(15), Article 11920. <https://doi.org/10.3390/su151511920>
- Gutiérrez-Páez, N. F., Santos, P., Hernández-Leo, D., & Carrió, M. (2023). A study of motivations, behavior, and contributions quality in online communities of teachers: A data analytics approach. *Computers & Education*, 201, Article 104829. <https://doi.org/10.1016/j.compedu.2023.104829>
- Hess, C., & Ostrom, E. (Eds.). (2007). *Understanding knowledge as a commons: From theory to practice*. MIT Press.
- Hornuf, L., & Jeworrek, S. (2023). The effect of community managers on online idea crowdsourcing activities. *Journal of the Association for Information Systems*, 24(1), 222–248. <https://doi.org/10.17705/1jais.00777>
- Hwang, S., & Foote, J. D. (2021). Why do people participate in small online communities? *Proceedings of the ACM on Human-Computer Interaction*, 5(CSCW2), Article 462, 1–25. <https://doi.org/10.1145/3479606>
- Madison, M. J., Frischmann, B. M., & Strandburg, K. J. (Eds.). (2014). *Governing knowledge commons*. Oxford University Press.
- Mustafa, S., Zhang, W., & Naveed, M. M. (2022). What motivates online community contributors to contribute consistently? A case study on Stackoverflow netizens. *Current Psychology*, 42(13), 10468–10481. <https://doi.org/10.1007/s12144-022-03307-4>
- Ostrom, E. (1990). *Governing the commons: The evolution of institutions for collective action*. Cambridge University Press.
- Ostrom, E. (2005). *Understanding institutional diversity*. Princeton University Press.
- Potts, J. (2018). Governing the innovation commons. *Journal of Institutional Economics*, 14(6), 1025–1047. <https://doi.org/10.1017/S1744137417000478>
- Potts, J. (2019). *Innovation commons: The origin of economic growth*. Oxford University Press.
- Redpill Linpro. (2023). *Att bygga gemensam digital infrastruktur för samverkan: Ett aktörsperspektiv*. Commons 2.0 project document.
- Sharp, A. R., Mpofu, N., Lankiewicz, E., Ajonye, B., Rambau, N. P., Dringus, S., Honermann, B., Erondy, N., Russell, A., Mwehonge, K., Aguiar, C., Ledan, N., & Kavanagh, M. M. (2024). Facilitators and barriers to community-led monitoring of health programs: Qualitative evidence

from the global implementation landscape. *PLOS Global Public Health*, 4(6), Article e0003293.
<https://doi.org/10.1371/journal.pgph.0003293>

Sweco. (2026). *A robust and agnostic architecture for future innovation and operations (RFI Commons 2.0)*. Commons 2.0 project document.

Wernerheim, T., Svensberg, L., Lindvall, J., Joelsson, E., Lundkvist, A., & Matthing, J. (2024). *Informationsstrukturer för lärande nätverk (ISLNd): Slutrapport för förstudie*. Karlstads kommun.