

Proceedings of the Paris Institute
for Advanced Study

Agentic AI in Data Commons and GLAM: What Role for the Model Context Protocol?

Casemajor , Nathalie ¹ Moreux, Jean-Philippe ²

¹ Institut national de la recherche scientifique, Canada

² Bibliothèque nationale de France

TO CITE

Casemajor , N., & Moreux, J.-P. (2026). Agentic AI in Data Commons and GLAM: What Role for the Model Context Protocol? In *Proceedings of the Paris Institute for Advanced Study* (Vol. 27). <https://paris.pias.science/article/agentic-ai-in-data-commons-and-glam-what-role-for-the-model-context-protocol>

PUBLICATION DATE

10/07/2026

ABSTRACT

The emergence of the agentic web, in which AI agents autonomously navigate online services to accomplish complex tasks, poses urgent questions for digital commons and GLAM institutions (Galleries, Libraries, Archives, and Museums). This paper reports on a cross-sectoral workshop held at the Bibliothèque nationale de France, which examined the Model Context Protocol (MCP) as a new infrastructure layer for regulating agentic access to open knowledge resources. Drawing on presentations by practitioners and researchers across the library, archive, museum, performing arts, open knowledge, and open science sectors, the paper documents emerging MCP experiments, identifies transversal governance challenges, and synthesizes the outcomes of a collective prospective scenario-building exercise. Key findings include: the reframing of MCP server deployment as an editorial act with significant implications for data

governance as a site of broader advocacy for the commons and public goods within the agentic web.

Contents

Acknowledgements	4
Introduction	4
Method: Knowledge Mobilization Workshop	5
The Agentic Web and Its Consequences for Knowledge Sharing: a Backend Turn	5
The Model Context Protocol (MCP): A New Infrastructure Layer for Agent Access	7
Uses and Challenges for GLAMS and Data Commons	9
Pioneering Experiments	10
Wikidata	11
French National Library (BnF)	12
Isidore	14
Future Directions	16
Prospective Scenario-Building: Projections Across Sectors	17
Libraries and Scientific Publishing	18
Archives	19
Museums	19
Performing Arts	20
Free Knowledge Organizations	21
Open Science	22
Discussion	23
From Invisibilization to the Shaping of Data Exposition	23
Use Cases	23
User Figures	24
Infrastructure and Development Capacity	24
Professional Competencies	24
Governance of the Agentic Web	24
Conclusion	25
AI Assistance Disclosure	25

Acknowledgements

This paper reports on a workshop organized by the authors as part of a research residency held by Nathalie Casemajor at the Paris Institute for Advanced Study, under the Distinguished Fellowship programme developed in partnership with the PostGenAI@Paris initiative (Sorbonne Université). The authors would like to thank all workshop participants for their valuable contributions to the discussions.

The workshop is part of a broader Franco-Québécois research project, *Digital Commons and AI*, conducted in collaboration with: (Québec, Canada) INRS Research Chair on Digital Commons, Centre de recherche interuniversitaire sur les humanités numériques (CRIHN / Interuniversity Research Centre on Digital Humanities), Bibliothèque et Archives nationales du Québec (BAnQ), Wikimedia Canada, University of Ottawa; (France): Laboratoire DICEN-IDF (Université Paris Nanterre), Wikimedia France, Laboratoire Théâtres & Médiations Numériques (TMNLab), Délégation générale à la langue française et aux langues de France (DGLFLF, General Delegation for the French Language and Languages of France), Ministry of Culture. With the support of the Consulate General of France in Québec and the Ministère des Relations internationales et de la Francophonie du Québec, through the Commission permanente de coopération franco-québécoise (CPCFQ / Permanent Commission for French-Québec Cooperation).

Introduction

What might the agentic web mean for digital commons and GLAMs (Galleries, Libraries, Archives and Museums)? The agentic web designates an emerging environment in which AI agents autonomously navigate the Web, interact with online services, and collect and synthesize information in order to accomplish complex tasks.

Against this backdrop of a broader paradigm shift in the Web, examining the implications for digital commons and heritage institutions becomes urgent. What infrastructures, interfaces, and regulatory frameworks will be required to absorb the anticipated surge in agent-generated traffic? How might open knowledge projects and heritage institutions contribute to shaping the standards structuring these new forms of interaction? And what synergies could be developed between public institutions and digital commons to collectively address these challenges? To explore these questions, we convened a cross-sectoral and transdisciplinary

workshop centered on the MCP protocol as a data representation interface, a data-mining tool, and a lever for regulating access to open resources.

Method: Knowledge Mobilization Workshop

The workshop was organized by the authors on June 12, 2026, at the Bibliothèque nationale de France (BnF), as part of a research residency at the Institut d'études avancées de Paris, with support from the PostGenAI@Paris Cluster (Sorbonne Université). It forms part of a broader Franco-Québécois collaborative project, led by the authors, examining digital commons and Web transformations driven by the expansion of AI systems.

Grounded in a partnership-based model of knowledge mobilization and transfer, the workshop convened 30 participants - professionals, researchers, and students - from a diverse range of sectors (see Table 1 in appendix): libraries, archives, museums, digital creation, government institutions, open knowledge nonprofits, data-sharing and research infrastructure networks, AI service providers, and research laboratories in AI, digital communication, heritage conservation, cultural studies, digital humanities, and film studies.


The workshop was organized around three successive moments: a framing session introducing the MCP protocol and its broader implications; case presentations illustrating concrete experimentation with MCP servers; and sector-based roundtables focused on prospective scenario-building and the identification of collective action strategies. The remainder of this paper documents the workshop presentations and key outcomes of the collective discussions, closing with a prospective reflection on research and action priorities for the next phase of the project.

The Agentic Web and Its Consequences for Knowledge Sharing: a Backend Turn

The opening presentation, by Raphaël Cousin (Research Engineer, SCAI – Sorbonne Université), established that LLMs are already exerting a profound impact on the web. Usage is migrating from websites to conversational assistants, a shift exemplified by the decline of Stack Overflow: monthly question volume plummeted following the launch of ChatGPT, accelerating the platform's decline ([Orosz, 2025](#)), while ChatGPT reached 900 million weekly active users by February 2026 (Figure 1). This structural shift in web usage patterns points to a critical juncture for data commons and open knowledge infrastructures: adapt to

the agentic web or risk the gradual obsolescence that has befallen Stack Overflow.

Figure 1. Evolution of ChatGPT Weekly Active Users. Source: esourcera




In practice, Cousin explains, the classic web journey involves querying a search engine and completing an action (e.g., looking up an artist and purchasing a museum ticket) through successive interfaces. In the agentic web, it gives way to a single request sent to an agentic chatbot (an LLM capable of calling and executing tools via an orchestrator), which handles all downstream actions through a unified interface.

This shift reorients attention from the frontend to the backend of the web (Figure 2): from what human users see (pages, buttons, browsers) to the program-structured interfaces through which agents operate: standardized HTTP requests (REST, SRU, SPARQL...) returning structured data (JSON, XML) rather than HTML pages. Backend standards already widely adopted in GLAM institutions – such as IIF (International Image Interoperability Framework, for structured access to digitized images and audiovisual resources), SRU (Search/Retrieve via URL, for querying bibliographic catalogues), and OAI-PMH (Open Archives Initiative Protocol for Metadata Harvesting, for metadata aggregation) – are cases in point.

APIs serve as the key mechanism for exposing these backend functions to external programs and, increasingly, to AI agents.

Figure 2. The Model Context Protocol (MCP) as a backend access layer for agents. Source: Raphaël Cousin.



Designing tools – a catalogue, for instance – suited to agentic chatbots thus means thinking in terms of function-call orchestration for accessing knowledge bases. Within this logic, what becomes critical is the quality of tool descriptions that LLMs can call upon. According to Cousin, it is this description, and this alone, that determines whether a model uses a tool effectively: "what the tool does, when to use it, what it returns". This reorientation carries significant governance implications: for GLAM institutions and data commons, regulating access to knowledge resources in the agentic web requires deliberate attention to backend infrastructure design: automated access architectures, computational rules, tool curation, and the crafting of tool descriptions – all dimensions that must be explicitly governed, rather than left to crystallize through market-driven defaults if public interest mandates are to be upheld.

The Model Context Protocol (MCP): A New Infrastructure Layer for Agent Access

A key architectural component of agentic systems is Retrieval-Augmented Generation (RAG), which at its core simply relies on a semantic database, as explained by Cousin. Rather than relying solely on knowledge encoded during training, a RAG system dynamically retrieves relevant information at inference¹ time. Documents are segmented into chunks, each transformed into a vector representation (an embedding)² that captures its semantic content. When a query is submitted, it is vectorized in turn, and the passages whose embeddings are closest in semantic space are retrieved – meaning that a query about a "castle"

1. Inference time refers to the moment when a trained model generates a response to a given input, as opposed to training time, when the model's parameters are learned from data.

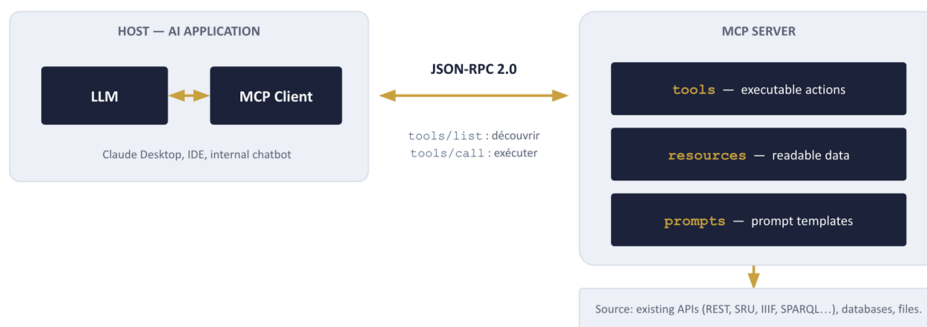
2. An embedding is a numerical vector representation of a document chunk in a high-dimensional space, where semantic similarity is captured as geometric proximity: chunks with related meanings are positioned closer together, enabling similarity-based retrieval.

may surface passages referring to a "medieval fortress," without any exact keyword match. These retrieved chunks are then injected into the LLM's context window, grounding the generated response in specific, retrievable knowledge.

Where RAG defines the internal logic of semantic retrieval at inference time, the Model Context Protocol (MCP) governs the standardized interface through which external tools and knowledge bases are exposed to and called by agents. MCP is an open standard introduced by Anthropic in late 2024. It has since been adopted by OpenAI, Google, and others, addressing a key limitation of standalone RAG systems: the absence of a common protocol enabling agents to dynamically discover and call heterogeneous tools and data sources at runtime.

As shown by Cousin in Figure 3, an MCP operates through two interconnected components. On one side, the *host* – the AI application (e.g., Claude Desktop, a specialized application, or an internal chatbot) – runs an LLM paired with an MCP client, one per connected server, responsible for discovering available tools and transmitting the model's calls via the JSON-RPC 2.0 protocol. On the other side, the MCP *server* exposes three types of resources to the agent: *tools* (executable actions), *resources* (readable data), and *prompts* (reusable query templates). An MCP server sits atop existing infrastructure, without requiring any modification to the underlying systems – APIs (REST, SRU, IIIF, or SPARQL), databases, or file systems.

Figure 3: MCP Components. Source: Raphaël Cousin



An MCP server functions not as a replacement for existing infrastructure, but as an orchestration layer unifying agent access across heterogeneous sources. As Cousin notes, MCP presents a clear scalability advantage: without it, every new agent connecting to every new data source requires a custom-built integration; with MCP, each data source exposes a single server, and every agent connects through a shared protocol. A single server can serve any agent – whether Claude, ChatGPT, or a custom in-house solution. For example, he explains, a natural language request such as "Find eighteenth-century maps of Paris and retrieve them in high resolution" can be handled by an MCP server exposing three functions: *search_documents* (SRU catalogue query), *retrieve_image* (IIIF high-resolution download), and *read_metadata* (OAI-PMH record retrieval). Crucially,

the agent chains these calls autonomously – querying the catalogue, then retrieving the images, then fetching the metadata – without any further human instruction. This autonomous sequencing is part of the agent loop.

For GLAM institutions and data commons, this architecture makes it possible to wrap existing resources (collections, metadata, and APIs) into MCP servers, making their knowledge bases directly accessible to many agentic systems³. In this context, well-described collections and robust APIs matter more than ever: they are a condition for remaining visible and accessible in the agentic web.

Uses and Challenges for GLAMS and Data Commons

What concrete applications might MCP enable for GLAM institutions and data commons? And what are the broader social and professional stakes of the agentic web for these organisations? For Julien Schuh (Prof. Université Paris Nanterre), MCP could do for knowledge infrastructures what IIF did for images: transforming isolated silos into shared, interoperable infrastructure. Where previous modes of access were limited to consultation or harvesting, MCP introduces an operationalization interface: a software layer through which institutions themselves can instrument access to their data, more flexibly than a full software stack would allow. Crucially, Schuh argues, MCP servers should not be understood merely as technical tools or applications: they constitute an editorial act. Deploying an MCP server is a new form of data editorialization: a deliberate decision about what is made accessible, how, and to whom. Understanding them as such is a prerequisite for any conscious and accountable engagement with them.

This editorial dimension carries significant implications for heritage institutions. RAG and MCP systems were designed for corporate environments (proprietary, normalized, homogeneous databases) whereas heritage repositories are fundamentally different: heterogeneous, historically layered, shaped by multiple and often divergent points of view. Schuh raises a series of questions that follow from this mismatch: what operations does an MCP server actually perform on a knowledge silo? What does it render visible, and what does it occlude? How does it interact with the limits of OCR (optical character recognition), and might semantic retrieval offer a pathway beyond them? And perhaps most consequentially: how are these operations described to the model? Since tool

3. At the time of writing, Mistral had not yet implemented support for the MCP protocol.

descriptions are written in natural language, the choice of words used to describe institutional functions will directly shape what is findable, and what is not.

Marta Severo (Prof. Université Paris Nanterre) broadened the discussion to situate MCP deployment within a wider set of responsibilities facing GLAM institutions. She opened with the question of *numérique responsable* (responsible and sustainable digital practices). Environmental sustainability is now part of France's regulatory landscape through the REEN framework (*Réduction de l'Empreinte Environnementale du Numérique*, 2021), which includes, for certain actors, an obligation to calculate their carbon footprint. Agentic systems compound this challenge: each additional loop and multiplication of automated requests adds to an already significant energy cost that organizations cannot ignore.

Beyond sustainability, Severo identified four interconnected dimensions of governance that GLAM institutions must navigate in the agentic web. The first is socio-political: as automation deepens, questions arise about where the human remains in the loop and what role human judgment retains. The second concerns data: the production of representative, high-quality datasets is a domain where GLAM institutions occupy a unique and critical position; they may be among the only actors capable of providing culturally and historically grounded context for training datasets and retrieval corpora that resist the homogenizing tendencies of commercially produced services. The third is governance proper: the choice of standards and software is inherently value-laden, and can be deliberately aligned with the public interest mandates of GLAM institutions. The fourth concerns security and trust: the protection of personal data and sensitive historical records, and the conditions under which users and communities can trust agentic systems, all warrant careful attention. GLAM institutions, as custodians of culturally and historically significant collections, are particularly concerned.

Pioneering Experiments

Several MCP experiments by public institutions and open knowledge organizations are already underway. In France, one of the first was the data.gouv.fr portal – the national open data platform – which launched its official MCP server in early 2026, enabling read-only search and exploration of open datasets via a public instance at mcp.data.gouv.fr. An extension of this server was proposed by Eudes Peyre, based on a dataset from the Ministry of Culture organized by topic. In this exploration, the topic "is no longer merely an editorial grouping: it becomes a machine entry point toward a bounded and contextualized

documentary space" (Peyre, 2026, p. 8). The experiment was conducted on an extension prototyped from a fork of the data.gouv.fr MCP server.

At the workshop, three further pioneering experiments were presented, covering libraries (BnF), open knowledge (Wikimedia), and academic research infrastructure (Huma-Num). Current experiments are limited to read-only access: the question of agents writing to or modifying resources has not yet been explored in the cases presented.

Wikidata

Philippe Saadé (AI/ML Project Manager, Wikimedia Deutschland) presented a case of MCP implementation on Wikidata, initiated in October 2025. Wikidata is a free, collaborative knowledge base maintained by Wikimedia Deutschland. Serving as a centralized data infrastructure for Wikipedia and other Wikimedia projects, it provides structured, multilingual data readable by both humans and machines, and currently contains over 122 million items released under a CC0 license (public domain).

Wikidata is structured around RDF triplets⁴, linking items to one another, to Wikipedia articles, and to external resources across the web. The MCP server exposes three complementary tools through the Wikidata API (Figure 4): graph traversal for navigating the network of linked items, SPARQL for precise structured queries, and a vector database for open-ended semantic search in natural language.

Figure 4: Conceptual schema of an MCP server for Wikidata. Source: Philippe Saadé.

4. RDF (Resource Description Framework) is a standard model for data interchange on the web; data is represented as triplets of the form subject-predicate-object ("Paris – is the capital of – France"), enabling machines to read and connect information across distributed sources.

Already part of the existing stack, SPARQL enables complex queries against Wikidata, but requires users to know in advance how the data is structured, which limits its accessibility. MCP helps bridge this gap by allowing agents to query Wikidata in natural language, without requiring prior knowledge of its underlying structure. A vector database was developed as a complementary tool, enabling semantic search across Wikidata. Where SPARQL allows for precise, structured queries, the vector database supports open-ended exploration in natural language.


Saadé highlighted a key design principle: rather than exposing a large number of tools that would risk confusing the LLM, the goal is to provide a short list of powerful, well-described tools capable of guiding users through a workflow. It is worth noting, however, that a workflow cannot be rigidly imposed through MCP.

French National Library (BnF)

Jean-Philippe Moreux (Head of the AI Mission, BnF) presented the BnF's preliminary experiment connecting Gallica resources to an MCP server. Gallica is the digital library of the BnF, providing free online access to over 10 million digitized documents from the BnF's collections and those of its partners (books, manuscripts, maps, photographs, newspapers, and audiovisual materials), most of which are digital reproductions of works in the public domain. The platform currently serves approximately 50,000 human users per day. For roughly two years, however, Gallica has been subjected to waves of scraping bots – and outright attacks – that threaten server stability and periodically block human access to resources. As of June 2026, bots account for 20-25% of the BnF's total bandwidth consumption.

A prototype MCP server has been connected to documentary databases (PDF documents, images) as well as descriptive metadata through Gallica's API (REST), providing access to both the resources themselves and the Gallica search engine (Figure 5).

Figure 5: Conceptual schema of an MCP server for Gallica resources (Bibliothèque nationale de France). Source: Jean-Philippe Moreux.



Moreux identified four potential benefits of MCP deployment in this context. The first is internal service improvement: an MCP server could serve as an interoperability layer across the BnF's dozens of databases, enabling the creation of new internal services. The planned overhaul of the Gallica search engine represents a timely opportunity to implement such an approach. Given the age and heterogeneity of the BnF's documentary infrastructure – Gallica is 25 years old, and the general catalogue spans several centuries – MCP offers a potential tool for making this legacy stack more interoperable.

The second is the regulation of access for AI model developers: MCP could provide a controlled channel through which commercial actors – including American and Chinese AI companies – are required to pass. Moreux framed this explicitly as a reactive measure, an attempt to regain some agency in a situation that institutions have largely been subjected to rather than shaped.

The third is the valorization of documentary expertise: MCP opens the possibility of embedding professional knowledge directly into the access layer. Just as a documentalist guides a user through a search, encoding domain expertise, professional rules, and experience into MCP tool descriptions could translate into concrete interventions – for instance, alerting a user that a query is poorly formulated. Schuh pushed the idea further, asking whether librarian personas could be embedded directly into the MCP layer as part of this reproduced documentary exploration chain.

The fourth is democratization of access to Gallica's resources. Moreux emphasized that MCP should not be thought of exclusively in terms of agents: it also concerns human users, both those with low digital literacy who require documentary mediation, and those arriving at the portal already equipped with agentic tools. However, MCP does not restore full control over the documentary exploration and interpretation chain: LLM-mediated access introduces an inherent opacity, whereby users interacting through conversational agents will receive different responses to identical queries depending on the model used. This also

fundamentally changes the nature of the interaction between a documentary web portal and its end users: where users once skimmed results and read selected excerpts directly, they now receive model-mediated summaries.

On the question of technical feasibility and maintenance, Moreux noted that MCP deployment does not alter the BnF's existing infrastructure, given the institution's long-standing API layer. MCP is not technically difficult to implement when robust API infrastructure and standardized access protocols are already in place. The key challenge, however, lies in capacity: opening an MCP server means absorbing a potential multiplication of requests, and the BnF's infrastructure bandwidth cannot be significantly expanded. The question of whether existing capacity can sustain the anticipated increase in usage remains open.

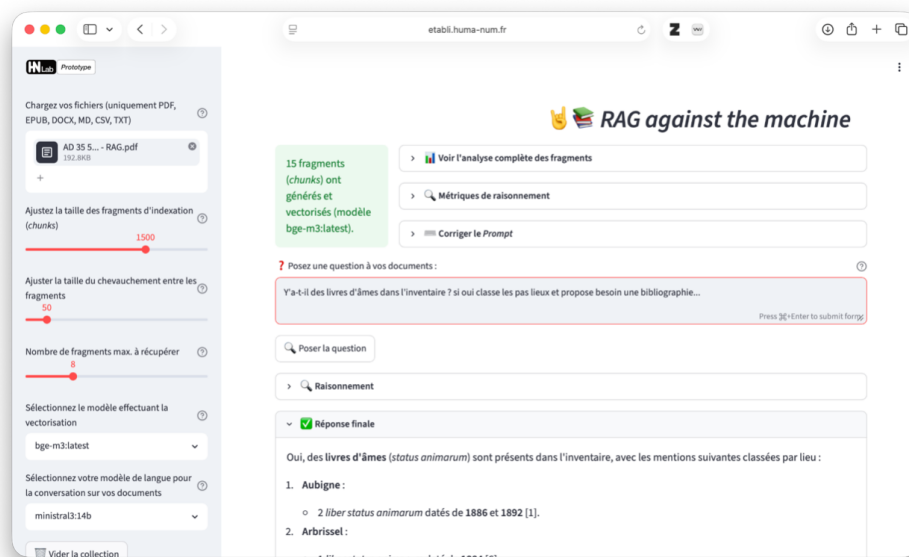
Isidore

Stéphane Pouyllau (Research Engineer, CNRS / Huma-Num) presented an MCP server experiment built around Isidore, an academic search engine for the social sciences and humanities (SSH). Isidore provides access to over 8 million documents: articles, archives, events, and more. The MCP project was developed in partnership with the company SPARNA, within the framework of Huma-Num, a digital research infrastructure serving SSH communities. It aligns with the FAIR principles (Findable, Accessible, Interoperable, Reusable), which promote open data and the development of collaborative tools.

The impetus for the experiment mirrors the situation described at the BnF: on the Isidore web portal, the balance between human and machine traffic has shifted dramatically. Approximately one year after the launch of ChatGPT, the human/machine ratio inverted; subsequent waves of indexing bots (targeting author pages in particular) have since compounded the problem. Some weeks see exclusively robotic traffic, and the API has been down for up to five out of six days. This raised a fundamental strategic question: should Isidore continue to maintain a frontend for human users, or become a platform designed exclusively for agents? The team's response was to pursue control rather than separation: managing human and agent traffic flows together rather than bifurcating the infrastructure.

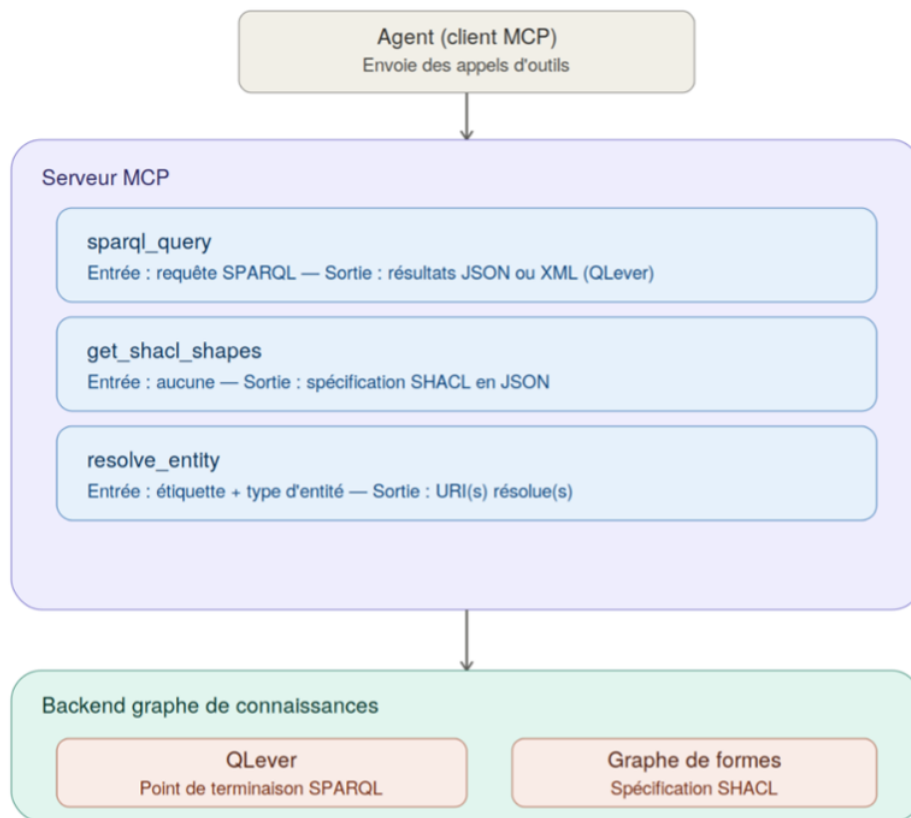
The MCP experiment, initiated in February 2026, was motivated by the need to respond to the growing demand from agentic chatbots. The objective was to build an MCP server capable of feeding a RAG and GraphRAG platform for both internal and external use (Figure 6), providing researchers with a tool to RAG over their own content, with pre-configured prompts.

Figure 6: Demonstration interface of an MCP server application built on Huma-Num resources. Source: Stéphane Pouyllau.



The MCP server builds on an existing web portal, an API, and a SPARQL endpoint. Pouyllau describes its organization around three core components (Figure 7): the first is a SPARQL service access tool, exposing the SPARQL endpoint via QLever (input: SPARQL query; output: SPARQL result in JSON or XML format). The second is a knowledge graph structure tool, exposing the SHACL specification of the knowledge graph (input: none; output: a JSON representation of the SHACL specification). The third is a named entity resolution tool, enabling agents to resolve named entities from their labels (input: label to search + entity type). The entire system runs on a lightweight virtual platform with modest resource requirements.

Figure 7: Conceptual schema of an MCP server for Isidore resources (Huma-Num LAB). Source: Stéphane Pouyllau.



Within a few weeks of implementation, over 7 million documents (through their enriched metadata) were made accessible. Significant effort was invested in mastering the exposition of the knowledge graph structure, including the possibility of exposing rules governing how the model should interpret and use the data. A named entity resolution tool completes the architecture.

However, the MCP server does not yet address the challenge of traffic surges: when a notable event triggers a spike in bot requests, the current approach still relies on redirecting traffic to a static page copy to contain the surge, rather than routing it through the live infrastructure.

Future Directions

These projects demonstrate that MCP adoption is underway among some GLAM institutions and open knowledge organizations, though all remain at an early stage. Two recommendations emerge from these experiments.

The first is strategic: for organizations with existing API infrastructure, establishing an official MCP presence is a means of ensuring visibility in a rapidly expanding ecosystem. Anyone can develop an MCP server in a matter of hours and connect it to an organization's publicly available API, with no need for authorization from the organization. In the BnF's case, several MCP servers providing access to the Gallica API already exist, developed not by the BnF itself

but by external actors⁵. This underscores the importance of developing official MCP servers and registering them in recognized MCP directories⁶. This allows both to ensure the quality and accuracy of the tool descriptions that will shape how agents access institutional resources, and to maintain institutional visibility and authority as the space is increasingly populated by third-party actors.

The second is methodological and concerns metrics: tools for evaluating the usage and effectiveness of MCP servers in documentary mediation contexts have yet to be developed. One avenue mentioned in the workshop discussion would be to design user-question benchmarks to assess how MCP servers respond to realistic queries. This would allow organizations to measure the quality of agent-mediated access to their resources. Peyre (2026) suggests that MCP experiments could be evaluated not only in terms of response quality, but also in terms of "the robustness of tool calls, resource selection, execution cost, and the capacity of different models to operate within a bounded documentary scope"⁷ (2026, p. 15). A further advantage of MCP, he argues, is its built-in traceability: the query chain preserves a record of every step, making errors visible in the response and correctable through editorial improvements to the catalogue.

Prospective Scenario-Building: Projections Across Sectors

The final part of the workshop took the form of a collective prospective scenario-building exercise. Participants were invited to imagine the impacts of the agentic web and the possible courses of action around MCP across the different culture and knowledge sectors represented at the workshop.

The exercise asked participants to project themselves five years into the future, to 2031, along two axes of projection (Figure 8): one representing the degree to which MCP becomes a structuring protocol for the agentic web (from highly structuring to rapidly obsolete), the other representing the degree of institutional uptake of MCP tools by GLAM and open knowledge organizations (from high uptake to low uptake). This framework generated four scenario configurations (Figure 8), defined by the intersection of MCP impact and institutional uptake: significant benefits, adverse consequences, marginal benefits, or limited consequences.

5. See, for example, Kryzo/mcp-bibliotheque_nationale_de_France.

6. See this official MCP registry: <https://registry.modelcontextprotocol.io/>

7. Translation by the authors.

Figure 8: Prospective Scenario Framework: MCP Impact and Institutional Uptake (2031 Horizon). Source: Nathalie Casemajor.



Participants worked in sector-based tables (museums, libraries, archives, performing arts, free knowledge, open science) identifying potential points of tension and opportunities for development within each scenario configuration.

Libraries and Scientific Publishing

The library sector table, documented by Julien Schuh, identified five areas of tension. The first one is existential: as machine-mediated consultation grows, the purpose of human-facing portals is called into question. Libraries observe stable on-site visits alongside a decline in physical document consultation, while their infrastructures face mounting bot traffic. The core challenge is preserving what lies at the heart of the library's mission – its documentary expertise and curatorial authority – as AI progressively displaces traditional access points to collections.

The second tension concerns data and metadata: digitization covers only a fraction of collections, and what remains undigitized risks becoming invisible to AI systems. Metadata quality remains a fundamentally human investment (AI can assist but not replace it), and the risk of AI-generated content feeding on its own hollow outputs compounds the problem. The third is legal: public data faces appropriation by large digital actors and the scientific publishing oligopoly. The position defended was that open licenses (CC-BY including commercial reuse) are the tool for keeping the field open to new entrants rather than consolidating dominant positions.

The fourth concerns institutional capacity: MCP remains largely unknown among library professionals, and users increasingly turn to AI rather than librarians, while some students request the right to opt out of AI use on ecological grounds. The fifth tension is governance and orchestration: existing resource-sharing mechanisms (such as CollEx) address optimization between libraries, not the

governance of reading and access practices. Who will orchestrate the MCP layer for the library sector remains an open question, with a real risk that this role defaults to private actors.

For the prospective exercise, the scenario envisions an MCP server built on open standards, providing access to research publications and data via persistent identifiers, developed by public research infrastructure actors (HAL, libraries, national infrastructures) in collaboration with the scientific community. The MCP would natively embed enriched metadata and trust markers readable by both humans and machines, under CC-BY licenses. Residual tensions include the control trade-offs of openness, the integrity threat posed by automated article generation, and the recognition that MCP remains, for now, a short-term tool for managing agent traffic rather than a long-term structural solution.

Archives

The archives table, facilitated by Marcel Bautista (AI Project Lead, BnF), focused on mapping the structural conditions and tensions relevant to MCP deployment in the archival context. The starting point was a shared diagnosis: heritage archives hold large deposits of metadata distributed across multiple silos, in heterogeneous and unstandardized formats. An overhaul of existing information systems is already needed, and participants noted the potential to align this with MCP-ready infrastructure development, subject to available resources.

Two MCP use cases were identified. The first is external-facing: a server designed to serve public users, with a noted risk of deepening the digital divide between users equipped to engage with agentic tools and those who are not. The second is internal: a server oriented toward professional workflows, enabling the reconciliation of internally siloed data and their linkage to external knowledge bases (Wikidata, BnF, other archival services). Both use cases presuppose the development of APIs that are currently lacking.

Several constraints were identified: legal obstacles around metadata that cannot be publicly disseminated; questions around technical realism and institutional capacity to project onto a fast-moving technological roadmap; resistance to agent personas and specialized agents; and the imperative to serve the full spectrum of users: academic researchers, amateur genealogists, and professionals engaged in legal or administrative research such as establishing inheritance rights or identifying rightful claimants.

Museums

The museum sector table, facilitated by Marta Severo, also worked from the observation that existing museum data is highly heterogeneous across institutions. Projected use cases remained primarily internal, even at the five-year horizon: participants found it difficult to project concrete external applications before the

internal layer is consolidated. On the external side, inter-institutional collaboration (developing shared networks between museums, with a designated bridging role) was identified as a more realistic near-term trajectory than direct public-facing MCP deployment.

Several governance dimensions emerged. The question of professional roles was central: participants expressed a clear will to invest in training as a condition for moving from prototype to production. The need to connect MCP development to existing mediation and communication professions was also underlined. Finally, participants identified the guarantee of data authority as a key issue: as agents increasingly mediate access to museum collections, institutions must maintain direct access services to their authoritative data, ensuring they remain the recognized source of reference rather than ceding that role to intermediary systems.

Performing Arts

The performing arts table, facilitated by Nathalie Casemajor (INRS), noted strong interest in MCP among several performing arts organizations, including the Réunion des Opéras de France. A distinctive challenge for the performing arts sector is event discoverability: because these organizations manage primarily event-based and time-sensitive data, their data structuring requirements differ substantially from those of heritage collections, with particular implications for audience forecasting.

Five areas of tension were identified. The first concerns information organization: data structures in the sector are not yet mature, archiving culture is less developed than in the heritage sector, and archival materials are often neither deposited, nor processed, nor digitized. The second concerns awareness and training: MCP remains largely unknown, and participants noted that the most effective communication approach would frame MCP not in technical terms but around the concrete risk of becoming invisible to AI agents. The third concerns acceptability: a desire to reaffirm human mediation coexists with concern about losing control over information production as LLMs become intermediaries, introducing errors and eroding user trust. The fourth is economic: institutions need clearer cost knowledge, investment rationalization tools, and metrics for measuring added value.

The fifth challenge concerns governance. Performing arts organizations navigate a structural tension between openness and siloed operations: while data sharing and interoperability would benefit the sector as a whole, competitive dynamics between theatres create significant friction around collaborative initiatives. At the same time, the question of institutional leadership remains unresolved: it is

unclear which actor would be best positioned to assume stewardship of MCP development and coordinate inter-institutional collaboration at the sectoral level.

For the prospective exercise, the table worked through the case of a small opera house. Structured event data and artwork-related data exist (including via the BnF), but cross-sectoral data models risk homogenizing institutional identities. Three questions structured the discussion: who internally takes ownership; whether to develop individual or mutualized servers (with a federated model of orchestrated mini-MCP servers proposed as a way to preserve local control while enabling coordination); and what the internal and external use cases are. Internal uses would include programming data, benchmarking, and trusted operational dashboards; external uses would center on event discoverability and new forms of audience mediation. On the economic front, the moment of website overhaul was identified as a strategic opportunity to integrate MCP development rather than treating it as an additional investment.

Free Knowledge Organizations

The discussion, documented by Stéphan-Eloïse Gras (Assistant Professor, CNAM-DICEN), centred on the Wikimedia movement as a case study, surfacing challenges that are at once infrastructural, epistemic, and organizational. The movement is navigating a conjuncture of declining readership and a perceived organizational lag relative to the pace of agentic AI development. A recurring question framed the discussion: if Wikipedia were being built today, what technologies would be used?

A central infrastructural concern is the invisibilization of Wikipedia within AI-mediated information environments. Participants questioned what mechanisms could guarantee that major AI providers would surface Wikipedia's MCP, and who holds legitimate authority to officialize and publish such servers. MCP governance was identified as inherently complex, involving the Wikimedia Foundation, local Wikimedia organizations, and partners such as Creative Commons, under significant resource constraints.

At the epistemic level, a fundamental tension structures the discussion: the opposition between synthetic content generation and Wikipedia's epistemic model, which has explicitly banned synthetic content. Participants identified a risk of epistemic friction, whereby MCP servers could function as closed circuits regulating access to documentary resources without guaranteeing output quality. The multilingual dimension also raised questions: participants asked whether the

MCP server called in a given interaction would correspond to the user's language of interaction, and how consistency across language editions could be ensured.

The table's discussion converged on a forward-looking framing: the priority for 2031 should be to ensure that a Wikimedia MCP ecosystem, ideally extended to open knowledge, open science, and open source partners, restores visibility to Wikimedia content within chatbots and AI agents, particularly for students and the general public. This raises the stakes of the patrimonial dimension: the challenge is not only to protect the content of an organization shaped by a past paradigm, but to define what a future-oriented Wikimedia knowledge infrastructure could look like.

Open Science

Discussion at this table, documented by Laurence Allard (MCF, IRCAV, Université de Lille), centered on the relationship between open science, open culture, and MCP, with discoverability and institutional capacity as organizing concerns. Four axes of tension were identified: availability (resources already accessible via catalogues, across formats and data types); licensing (a key friction point, as scraping practices routinely disregard license conditions); governance (the risk of institutional disengagement if ordinary interfaces are not developed); and institutional capacity (the role of documentary institutions in making their resources visible within AI-mediated environments).

The second part of the discussion took a more projective turn. Participants articulated a vision of MCP not as a passive conduit but as a politically agentic infrastructure. A desirable MCP within open knowledge ecosystems would operate across four dimensions: accessibility, guaranteeing equal access to open and non-commercial content against the barriers introduced by commercial LLM subscriptions; transparency, surfacing the gendered and ethno-racial biases of commercial LLMs; eco-conditionality, flagging redundant queries and redirecting them toward classical search engines or human intermediaries such as librarians; and protection of openness, prioritizing open licenses and reciprocity-based frameworks (Creative Commons, UseRight) with opt-in mechanisms for researchers.

The scenario converged on a political horizon: MCP development governed by ambitious open knowledge policies, protecting open epistemic infrastructures against commercial and ecocidal forms of predation.

Discussion

The presentations and exchanges held during the workshop brought out a number of transversal issues regarding MCP's role as a regulatory layer for agentic access, as a mechanism for data representation and contextualization, and as a vector of documentary mediation for non-expert users.

From Invisibilization to the Shaping of Data Exposition

The most consistently shared concern across tables was the risk of invisibilization: libraries, Wikipedia, performing arts organizations, and GLAM institutions alike identified the displacement of their resources within AI-mediated information environments as a central threat, whether as a loss of curatorial authority or a failure of discoverability. As shifts in online behaviour are already visible, the MCP protocol offers institutions a concrete lever to respond.

Creating an MCP server was characterized as an editorial act: it involves exposing a knowledge graph and specifying the rules governing how models should explore the data. Control over this interface layer was identified as a new site of value creation for heritage institutions and knowledge commons. For a long time, value resided in data itself; it now increasingly resides in how data is made accessible. This aligns with Peyre's observations, which distinguish two complementary dimensions of a shared space of trust for the agentic web: "documentary governance, focused on scope, sources, schemas, and data traceability; and orchestration governance, focused on the conditions of execution and tool selection" ([Peyre, 2026](#)). MCP thus represents a central negotiation space for mediating access to the past, at a moment when such mediation will occur massively and directly via LLMs.

Use Cases

Two main categories of use cases were identified across tables. Internal uses centre on the orchestration of database query functions and the interoperability of currently siloed systems. Museums and archives project primarily internal use cases at a five-year horizon, while libraries, open science and open knowledge organizations are already engaging with external-facing implications. External uses centre on documentary mediation toward end users: transmitting professional documentary expertise and democratizing access to resources, including as an alternative to complex advanced query forms.

User Figures

Participants cautioned against developing MCP exclusively for machine access at the expense of human access. The opposition between human and machine access was identified as a false dichotomy: it is necessary to also account for equipped human users who arrive at portals via agents acting on their behalf. At the same time, participants noted the risk of presupposing that all users are equally equipped: while basic access to agentic chatbots is often free, sustained use requires a commercial subscription, inserting a market intermediary between users and the open knowledge resources they seek to reach.

Infrastructure and Development Capacity

MCP can interface with existing systems, and development costs may be relatively low. However, deployment can generate significant access volumes, raising bandwidth constraints and contributing to the broader growth of agentic traffic. A structural gap was identified between organizations that have already developed APIs and can add this interface layer, and those that lack the foundational infrastructure to do so. The maturity of existing data infrastructure varies considerably across sectors. Some institutions present at the workshop hesitate even to publish new metadata, given insufficient resources to handle the resulting demand for access to source documents.

Professional Competencies

A shared concern across tables was that MCP remains largely unknown among professionals across all sectors, and resources are scarce. Participants identified new professional competencies required for MCP development, particularly around tool description. Questions of digital literacy were raised regarding how MCP integration relates to broader professional roles and more transversal skill sets. The importance of making MCP servers auditable, including their biases, was also noted.

Governance of the Agentic Web

Governance emerged as a structuring question across all tables: who develops, publishes, and maintains MCP servers, and under what authority. A shared concern was that, by default, this role will fall to private actors. Beyond institutional governance, participants raised the governance of the protocol itself. MCP has already evolved through several iterations and will continue to develop. How could it evolve to be adapted to GLAM and data commons uses? RAG and MCP systems were designed for corporate environments, with proprietary, normalized, and homogeneous databases, whereas heritage repositories present fundamentally different conditions. Participants underlined the importance of

communicating these stakes to GLAM communities, and of contributing to the design of the protocol itself as a form of governance by design, oriented toward open governance of the agentic web. This is the thrust of Alek Tarkowski's (Open Future Foundation) call for Wikimedia to be granted "a seat at the table" within the Agentic AI Foundation – whose mission is to govern the open protocols and frameworks underpinning AI development – in order to represent the interests of the broad knowledge commons ([Tarkowski, 2025](#)). This argument could be extended to GLAM networks such as Europeana, suggesting a strategic alliance between knowledge commons and cultural public goods in defending their shared vision and developmental needs for the future of the agentic web.

Conclusion

As a new site of value creation and regulation, the MCP protocol offers institutions an easy-to-deploy and concrete lever for responding to the growth of agentic access. Discussions pointed to the importance of GLAM institutions and open knowledge organizations positioning themselves rapidly. However, MCP deployment may reinforce existing institutional capacity gaps, as it relies on existing API infrastructures and requires the development of new professional competencies. Another risk lies in focusing primarily on agentic users, whereas the development of machine interfaces should not come at the expense of non-equipped human access interfaces. Finally, the governance of the MCP protocol may serve as a broader site of advocacy for greater representation of the commons and public goods within the governance of the agentic web.

AI Assistance Disclosure

The authors acknowledge the use of Claude Sonnet 4.6 (Anthropic) as a writing assistance tool for translating portions of this text from French to English and for rephrasing select passages. The author retains full responsibility for the final content, analysis, and interpretations presented in this article.

Bibliography

- Orosz, G. (2025). The Pulse #134: Stack Overflow Is Almost Dead. *The Pragmatic Engineer*.
<https://newsletter.pragmaticengineer.com/p/the-pulse-134>.
- Peyre, E. (2026). *Gouvernance et outils conversationnels génératifs: Expérimentation d'un espace de confiance de la donnée*. <https://hal.science/hal-05655006>.
- Tarkowski, A. (2025). *Why Wikimedia needs a seat at the Agentic AI Foundation*. Open Future.
<https://openfuture.eu/blog/why-wikimedia-needs-a-seat-at-the-agentic-ai-foundation/>.