

Museum Monitoring: an environmental monitoring dataspace using The Things Network, Solid, and LDES

Arthur Vercruyse¹, Ben De Meester¹, Julián Rojas¹ and Dieter Suls²

¹IDLab, Department of Electronics and Information Systems, Ghent University - imec, Belgium

²Fashion Museum Antwerp, Nationalestraat 28, 2000 Antwerpen, Belgium

Abstract

Museums increasingly deploy environmental monitoring systems (e.g., data loggers and sensor networks) to support preservation of cultural heritage objects. However, monitoring data is often captured and accessed through proprietary vendor software, which makes reuse, integration, and controlled cross-institution sharing difficult—particularly in object loan scenarios. This paper reports on MuMo (Museum Monitoring), a three-year applied research project that explored how dataspace principles can be applied to environmental monitoring in real museum settings. Rather than replacing existing systems, MuMo extends a legacy monitoring dashboard with semantic data modeling, Linked Data Event Streams, and Solid-compliant data access management. We present the system design, where environmental monitoring data is semantically described and published with access controlled fragments as a stream. The semantic links between these fragments allow clients to prune irrelevant branches while providers can restrict access at natural boundaries (e.g., per location or sensor). Further, we describe how MuMo is used in practice through a set of in-use scenarios, where the system enables (1) decentralized data publication for longitudinal analysis of environmental conditions, (2) selective cross-institutional data sharing during object loans, and (3) client-side data integration and aggregation across independent deployments. In practice, this dataspace-oriented deployment reveals design trade-offs while preserving institutional autonomy: museums control their data and authorization policies end-to-end, supporting trusted data sharing across organizational boundaries. The deployed system and its insights are thus relevant to a broad range of (Digital Humanities) projects that involve long-running data integration under distributed governance.

Keywords

Digital Humanities, Dataspaces, Solid, LDES, Museum Monitoring

1. Introduction & Motivation

In cultural heritage institutions, environmental parameters such as temperature, relative humidity, and light exposure directly influence the long-term preservation of collection objects [12]. Museums therefore invest in digital infrastructures for continuous environmental monitoring and traceable reporting [12, 14].

A wide range of sensing technologies is already used in practice, including data loggers and long-running wireless sensor network deployments in museum buildings [15]. Yet these systems typically remain tightly coupled to proprietary dashboards or local infrastructures [1]. As a consequence, monitoring data is frequently siloed: it is difficult to combine measurements across installations, to align monitoring data with other institutional sources, or to selectively share a well-scoped subset of measurements with external partners under enforceable access control. This fragmentation makes cross-organizational reuse unnecessarily costly.

This paper reports on MuMo (Museum Monitoring)¹, a three-year applied research project², funded by the Flemish Government, that investigates how museum environmental monitoring can be made more reusable, with a focus to not force museums to replace their day-to-day operational tooling. The

SemDH 2026: Third International Workshop of Semantic Digital Humanities. Co-located with ESWC 2026, May 10, 2026, Dubrovnik, Croatia

*Corresponding author.

†These authors contributed equally.



© 2026 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<https://faro.be/project/museum-monitoring-tool>

²<https://www.vlaanderen.be/cjm/nl/cultuur/cultureel-erfgoed/subsidies-cultureel-erfgoed/projects/subsidies-cultureel-erfgoed/internationale-landelijke-en-bovenlokale-cultureel-erfgoedprojecten>

project included both hardware/firmware and a software dataspace layer; this paper focuses on the software and dataspace aspects.

MuMo is motivated by the following monitoring needs:

1. **Operational oversight**

Staff need near-continuous insight into readings, time-series views, and alerts when conditions leave acceptable ranges.

2. **Long-term documentation**

Museums need access to historical exposure conditions over extended periods, so they can reconstruct how an object's environment evolved across locations and organizations, and throughout its lifecycle.

3. **Selective collaboration and sharing**

When multiple parties are involved—most notably during loans—access must be manageable and bounded [11]. Practically, this means permissions that can be limited by (i) a subset of sensors/rooms, (ii) a defined time window, and (iii) the involved organizations.

The loan scenario is a key stress test. Loan agreements frequently include environmental constraints³, and the lending institution needs trustworthy insight into the conditions experienced by an object while it is hosted elsewhere. At the same time, the borrowing institution is typically unable—and often unwilling—to expose its full internal monitoring landscape. In practice, this leads to ad-hoc exports and manual reporting, introducing delays, duplicate effort, and reduced transparency [9, 8].

To make the discussion concrete, we use a running example throughout this paper: **Museum A lends a painting to Museum B for a temporary exhibition**. Museum A needs to verify that environmental conditions at Museum B remain within the agreed conservation thresholds during the loan period. At the same time, Museum B should not have to expose its full monitoring infrastructure. This scenario illustrates each of MuMo's design requirements—operational oversight, long-term documentation, and selective sharing—and will be revisited as we present the data model, publication layer, governance, and in-use scenarios.

Concretely, this paper addresses three research questions:

1. **RQ1** How can dataspace principles be applied to extend legacy museum monitoring systems for interoperable, semantically described data publication?
2. **RQ2** How can selective, revocable cross-institutional data sharing be realized without centralizing governance or identity management?
3. **RQ3** What practical lessons emerge from deploying such an architecture in operational museum settings?

In addressing these questions, we make three contributions: (1) an in-use account of deploying dataspace principles for environmental monitoring in museums, (2) a concrete architectural integration of semantic modeling, event-based publication, and decentralized governance interacting with legacy systems, and (3) empirically grounded lessons on aligning access control and data sharing mechanisms with institutional practice.

After describing the project context in Section 2, we present MuMo's architecture and approach, and implementation in Sections 3 and 4, respectively, and illustrate its use in practice through in-use scenarios in Section 5. We then discuss lessons learned in Section 6, reflect on its implications for the Digital Humanities community in Section 7, and position our work with respect to related work in Section 8. We conclude in Section 9.

2. Background and Design Constraints

To avoid ambiguity, we use MuMo v1 for the earlier prototype and MuMo v2 for the work presented in this paper. Unless stated otherwise, “MuMo” refers to MuMo v2.

³<https://collectionstrust.org.uk/wp-content/uploads/2017/11/Loans-out-lending-objects.pdf>, p11

MuMo v1 originated from the observation that environmental monitoring data in museums is difficult to access and reuse outside vendor dashboards. Monitoring equipment is often expensive and relies on proprietary, black-box devices that produce data locked to a single vendor platform—neither interoperable nor easily exchanged with other systems, and therefore underused for strategic purposes. To demonstrate that this could be done differently, the Fashion Museum Antwerp built an initial in-house prototype in 2019 using LoRaWAN shields, an Arduino, and a Raspberry Pi. Based on the prototype’s demonstrated potential, the museum and the University of Antwerp applied for funding to extend the work.

The project evolved through several phases: an initial prototype phase (2019), in which the Fashion Museum Antwerp identified the problem and built a first working prototype; a funded PoC phase (MuMo v1, 2020–2021), in which MoMu and the University of Antwerp developed a proof-of-concept monitoring setup using open-source hardware and a PHP dashboard under an ‘Innovative Project’ grant from the Flemish Government; and a funded development phase (MuMo v2, 2022–2025), a three-phase ‘Cultural Heritage Project’ also funded by the Flemish Government, in which the system was extended with custom hardware and a dataspace-oriented software layer before being put into operational use in 2026. Figure 1 gives an overview of these phases.

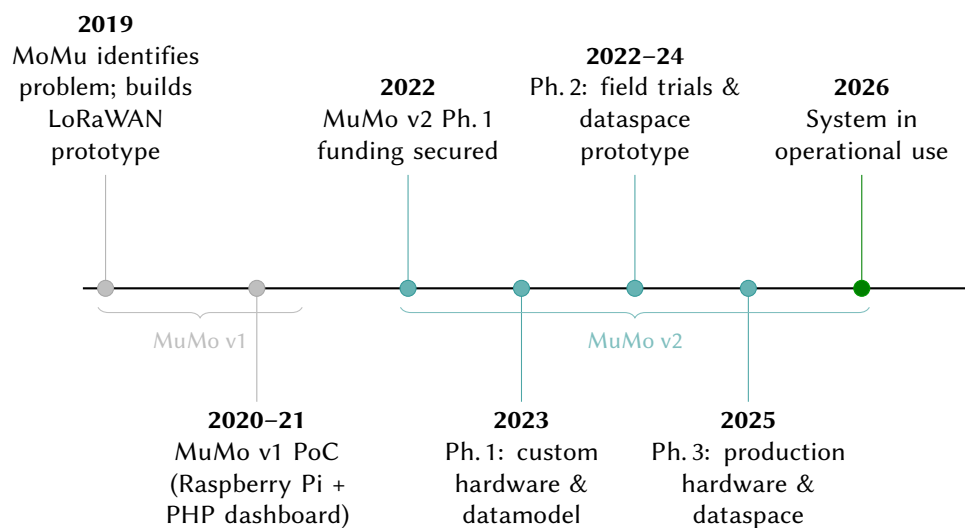


Figure 1: MuMo project timeline.

2.1. Starting point: MuMo v1

MuMo v2 starts from MuMo v1, which already supported core operational monitoring: ingestion and persistence of sensor readings, time-series visualization, alerting, and basic export functionality. MuMo v1’s dashboard also supported user and group management aligned with museum workflows, enabling staff to manage access within an institution in terms of responsibilities and spaces. In practice, these groups reflect a nested location hierarchy: from the museum or site level down to individual rooms and storage areas, and—when needed—even to fine-grained containers such as cabinets, shelves, or specific boxes.

At the same time, MuMo v1 made visible the limitations that motivated MuMo v2’s focus: data remained largely bound to a single dashboard instance, exports were the primary sharing mechanism, and interoperability and cross-institution governance were not first-class concerns.

2.2. Scope: extending rather than replacing existing systems

MuMo v2 is explicitly not a “rip-and-replace” effort. Museums already depend on established dashboards and workflows that are difficult to change in day-to-day practice. Consequently, MuMo v2 keeps a

dashboard-centric workflow as the primary operational interface for staff—where sensors are configured, readings are inspected, and alerts are handled—while adding a complementary data layer aimed at long-term reuse and cross-institution sharing.

This design stance reflects practical constraints in museum IT: successful changes must remain compatible with institutional autonomy, limited technical staffing, and long-running deployments. MuMo v2 therefore preserves the operational loop (“measure → view → react”) while enabling monitoring data to be reused and selectively shared beyond a single dashboard when needed.

2.3. What MuMo v2 adds: new hardware and a dataspace-oriented layer

Building on this baseline, MuMo v2 upgrades both the sensing infrastructure and the data layer. On the hardware side, it replaces the Raspberry Pi/Wi-Fi setup with custom low-power sensors designed for long battery life (at least x months on a single charge) and LoRaWAN-based uplink communication⁴. LoRaWAN is used for device-to-gateway transmission (low power, limited payload), after which a powered gateway bridges measurements to the internet.

Measurements captured by the MuMo v2 devices are transmitted to off-the-shelf gateways and routed through The Things Network before being ingested into the existing monitoring dashboard. On the software side, MuMo v2 adds the capabilities needed for reuse and selective sharing: (i) semantic representation of monitoring data, (ii) incremental/event-based publication, and (iii) an authorization approach that can operate across institutional boundaries while remaining administratively feasible for museum staff.

Overall, MuMo v2 is shaped by the need to support long-running monitoring data while enabling selective, revocable sharing across organizations—without disrupting the existing operational setup.

3. System and Approach Overview

In this chapter, we discuss the developed system that (1) keeps operational oversight, aligned with the familiar day-to-day monitoring workflow by extending the existing dashboard application; (2) provides long-term documentation through a semantic data model and representation (Section 3.1); and (3) enables selective data sharing through group-and-time-constrained access control (Section 3.3).

We further detail the governance layer and operationalization in Sections 3.3 and 4, respectively.

3.1. Data model and semantic representation

Semantic Web technologies are central to MuMo’s approach: they provide the interoperability layer that enables data to be published, discovered, and integrated across independent deployments without requiring shared infrastructure or coordinated schemas. MuMo provides a semantic representation—using RDF as the underlying data model—of all information managed in the dashboard, both environmental observations (e.g., temperature, humidity, light exposure) and sensor configuration metadata (e.g., group/location). By grounding all published data in established Linked Data standards, MuMo ensures that any standards-compliant client can consume its data without bespoke integration effort.

3.1.1. Modeling sensors and observations

MuMo adopts the Flemish OSLO (Open Standaarden voor Linkende Organisaties) vocabularies and application profiles [2]. These vocabularies and application profiles—maintained by the Flemish governmental organization Digitaal Vlaanderen—reuse established international semantic standards for interoperable cross-organizational data exchange.

In particular, MuMo uses the OSLO Sensoren en Bemonstering application profile⁵, which reuses the W3C Semantic Sensor Network ontology (SSN/SOSA) [5, 10]. SSN/SOSA is particularly suitable

⁴<https://www.thethingsnetwork.org/docs/lorawan/architecture/>

⁵<https://data.vlaanderen.be/doc/applicatieprofiel/sensoren-en-bemonstering/>

because it separates (i) the sensor, (ii) the observation, and (iii) the feature of interest being observed⁶. This allows MuMo to represent sensor redeployment and contextual change (e.g., a sensor moved to a different space) without rewriting historical observations that were produced under earlier conditions. In our running example, this means Museum B can deploy a sensor near the borrowed painting and later reassign it—without losing the observation history during the loan period.

3.1.2. Representing changing context as configuration events

Monitoring data exhibits two different dynamics:

- Observations evolve continuously and are typically appended over time.
- Configuration context evolves discretely when sensors are moved or reconfigured.

MuMo makes this distinction explicit by publishing configuration metadata as a versioned event stream, allowing consumers to reconstruct which context applied when an observation was produced.

Information about the group of the sensor (the location) is published with the versioned sensor metadata to make sure it is fully expressed in terms that the museum staff are already using (groups representing rooms, objects, loan packages, etc.).

3.2. Publication layer: Linked Data Event Streams

To publish this append-only log of measurements, MuMo applies the Linked Data Event Streams (LDES) technical standard⁷. LDES is maintained by SEMIC and actively promoted in Flanders (Digitaal Vlaanderen) for interoperable data sharing [6, 19].

LDES specifies how evolving datasets can be published incrementally in linked interoperable data fragments. Applying LDES supports any MuMo client application to automatically ingest incremental events across organizational boundaries without relying on centralized query endpoints, and as such allows two key consumption modes for monitoring data: replication of historical data and synchronization with newly produced data. We introduce the LDES principles, after which we clarify how we applied LDES in MuMo.

3.2.1. LDES Principles

In LDES, data fragments are typically organized as a tree: measurements come in as events, these events are partitioned into published resources (called *fragments*), and each fragment exposes typed links to other fragments through machine-interpretable relations in RDF. Starting from an *LDES entrypoint*, these typed links allow automatic traversal of the fragment tree. Listing 1 illustrates this with a partitioning by sensor: the entrypoint fragment links, via `tree:EqualToRelation` on `sosa:madeBySensor`, to a dedicated fragment for each sensor. In practice, such partitionings can be composed: once a client has navigated to the fragment for the relevant sensor, that fragment can in turn act as the root of a second tree (e.g., partitioned by time windows such as month/day), enabling clients to narrow down first by source and then by interval without scanning unrelated observations.

3.2.2. Fragmentation as a design lever

A central design choice in LDES is the fragmentation strategy: how events are grouped into fragments. MuMo uses fragmentation not only as a performance mechanism, but as a way to reflect how museums work with monitoring data and how sharing is scoped in practice.

In MuMo, observations are fragmented along the operational dimensions that conservator and technicians naturally use: group, sensor, and time. Concretely, observations are published under a hierarchy of the form:

⁶A primer in Dutch can be found here <https://museummonitoring.github.io/MUMO-Primer/>

⁷<https://semiceu.github.io/LinkedDataEventStreams/>

```

@base <https://mumo.faro.be/data/> .
@prefix tree: <https://w3id.org/tree#> .
@prefix ldes: <https://w3id.org/ldes#> .
@prefix sosa: <http://www.w3.org/ns/sosa/>.

# LDES endpoint
<stream>
  a ldes:EventStream ;
  tree:view <by-sensor> .

# <by-sensor> fragment
<by-sensor>
  a tree:Node ;
  tree:relation [
    a tree:EqualToRelation ;
    tree:path sosa:madeBySensor ;
    tree:value <https://mumo.faro.be/sensors/sensor-1> ;
    tree:node <by-sensor/sensor-1> ;
  ] ;
  tree:relation [
    a tree:EqualToRelation ;
    tree:path sosa:madeBySensor;
    tree:value <https://mumo.faro.be/sensors/sensor-2> ;
    tree:node <by-sensor/sensor-2> ;
  ] .

```

Listing 1: LDES fragmentation overview example, first fragmenting on location, then on group and lastly on time.

group-x / sensor-x / year / month / day

This keeps fragments small and stable over time (e.g., daily fragments do not grow indefinitely), while retaining a navigable structure that allows clients to focus on the relevant group, sensors, and time window (example shown in Figure 2). In the loan example, Museum B publishes the borrowed painting’s monitoring data under a dedicated loan group; Museum A’s client can then navigate directly to that group’s subtree and retrieve only the relevant observations. Note that the group-level fragmentation is flattened: the fragment tree does not encode the group hierarchy itself. Instead, parent–child relationships between groups are provided separately in the group metadata, which is sufficient for navigation.

This publication model suits cross-institutional settings because it allows consumers to process only the data they are authorized to access, without requiring providers to expose tailored query endpoints.

3.3. Governance layer: Solid-based decentralized sharing

Cross-institution collaboration (especially loans) requires selective, revocable access across organizational boundaries. Meanwhile, museums want to keep operational control over their monitoring infrastructures.

MuMo addresses this—conform to and supported by related work [16]—by combining decentralized LDES publication with Solid-based authorization. Solid provides a framework in which storage (called *Solid Pods*), identity, and access control are decoupled from specific applications, enabling the same published resources to be reused by multiple clients without funneling all access through a centralized service [3, 13, 20].

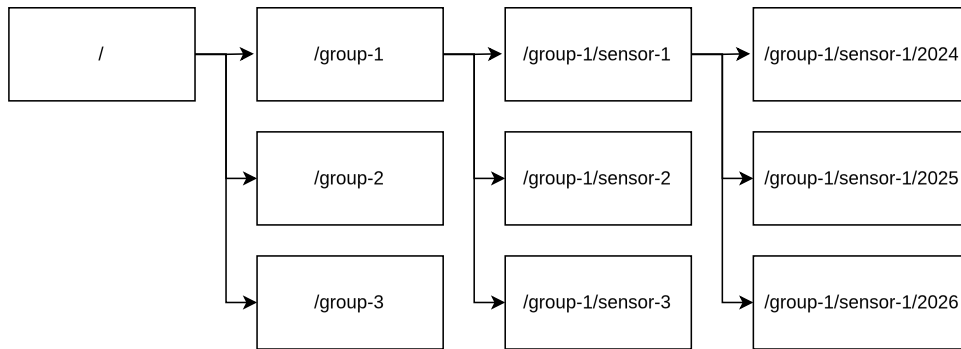


Figure 2: LDES fragmentation overview example, first fragmenting on location, then on group and lastly on time.

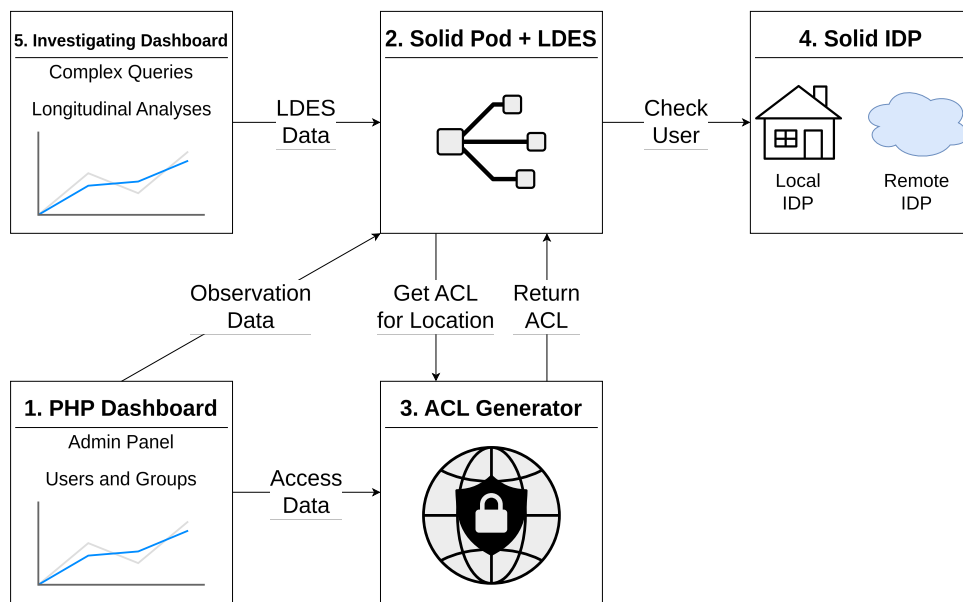


Figure 3: MuMo deployment overview with four components, an Identity Provider, a dashboard, a solid pod with the LDES publications, and an ACL file generator.

MuMo’s governance layer consists of (1) institutional endpoints with (2) group-based authorization (3) aligned with the LDES fragmentation strategy.

3.3.1. Institutional endpoints

In MuMo, Solid Pods act as institutional data endpoints rather than user-centric storage. The practical effect is that a museum can publish monitoring resources as long-lived datasets under its governance, while granting external parties access only to the parts required for a collaboration (e.g., a loan group) and revoking that access afterward.

3.3.2. Authorization model: group-based sharing as an operational advantage

In Mumo, the authorization model is group-based: access is granted at the level of sensor groups (mirrored from the groups managed in the dashboard) rather than individual observations. This is sufficient in museum practice: loan scenarios typically require a partner to access the monitoring data relevant to a particular object’s location over a time period, while keeping all other monitoring data private.

3.3.3. Policy-aligned fragmentation (LDES + Solid together)

Solid authorization mechanisms such as Web Access Control (WAC) enforce access rules over resources and their representations [3, 4].

Because LDES allows the publisher to decide how events are partitioned into resources (fragments), MuMo chooses a fragmentation strategy that makes the group the first-class boundary: the group subtree becomes the unit of sharing.

In the deployed system, authentication is enforced at the first level of the fragment tree: users are granted access to the subtree for a given group, which implicitly includes all sensors and time slices published under that group. The corresponding ACL files are generated from the dashboard configuration; granting a user access to a group also grants access to all of that group's descendants. Because the fragment tree uses a flattened group partition, a single change in the dashboard can therefore update the effective authorization for many fragments at once.

The same structure also makes it explicit that finer-grained authorization would be technically possible (e.g., sensor-level or time-range-level restrictions), but MuMo does not activate those options to remain aligned with what museum staff can practically configure in the dashboard. This only goes as far as the fragmentation allows, for example you cannot give access to half a day of data, as a half day is not a fragment, only full days exist. In the current deploy an ACL file exists for the first layer in Figure 2, but each entry point to a subtree could have an associated ACL file.

4. Implementation

Figure 3 shows how MuMo bridges the dashboard with synchronized Solid-based identity and authorization at runtime.

4.1. Runtime components

MuMo consists of four main components.

Component 1 – Solid Identity Provider (IdP).

The IdP is where users log in (with Solid-OIDC). After login it identifies them by their WebID, which MuMo uses as the handle for access control. A deployment can run its own IdP to issue WebIDs locally, but it can also accept WebIDs from other providers, so identity does not have to be centralized.

Component 2 – dashboard (admin & group management).

The dashboard is the control plane where administrators manage groups and group hierarchies (including recursive nesting), user-to-group assignments, and cross-institution sharing by granting group access to WebIDs. WebIDs may originate from the deployment's local Solid IdP or an external provider such as Inrupt; in the dashboard they function as stable identifiers for granting/revoking group access.

Component 3 – Solid Pod hosting LDES resources (data plane).

The Solid Pod hosts the published resources (as the LDES fragment tree) and enforces WAC authorization before serving protected resources. Requests are interpreted in terms of which protected subtree (group/sensor/time path) is being accessed and whether the requester's authenticated WebID is authorized. Data for these resources originate from the dashboard.

Component 4 – ACL generator (policy materialization service).

The ACL generator bridges the dashboard's group model with Solid's resource-level authorization. It receives synchronized configuration (WebID → groups and the group hierarchy) and materializes WAC ACL documents on demand. Operationally, its function is: given a request targeting a particular group, return the effective ACL representation for that scope. This avoids manual ACL maintenance and keeps the Solid-side authorization view consistent with the dashboard configuration.

Component 5 – Investigating dashboard.

A additional client-side dashboard—published at <https://museummonitoring.github.io/graphs/>—retrieves sensor descriptions first, determines the user's authorized scope, and then incrementally consumes only the relevant observation streams. This allows consumers to combine data from multiple sources in a

unified user experience, without requiring centralized aggregation. This is consistent with the principle that integration occurs at the point of use.

4.2. Flows

Control path:

1. Admin configuration: an administrator updates group membership in the dashboard by associating WebIDs with groups and maintaining the group hierarchy.
2. Sync to policy service: these settings are synchronized to the ACL generator as (WebID → groups) plus the group hierarchy.

Data path:

1. Data update: a measurement is ingested in the dashboard.
2. User data request (data path): a client requests an LDES resource from the Solid Pod.
3. ACL resolution: before serving the resource, the Solid Pod determines which ACL scope applies (which group subtree is being accessed) and consults the ACL generator for the effective ACL.
4. Authentication + authorization: the Solid Pod validates identity via the relevant IdP (local or remote) and evaluates whether the authenticated WebID is granted access by the generated ACL.
5. Response: if authorized, the Solid Pod serves the requested LDES resource; otherwise it denies access.

5. In-Use Scenarios

This section illustrates how the MuMo dataspace architecture is used in practice by museum professionals to analyze, share, and contextualize environmental monitoring data.

5.1. Scenario 1: Analyzing Environmental Conditions Over Time

The primary use of the advanced dashboard is to support museum staff in assessing the long-term “health” of collection objects by analyzing environmental conditions over time. Users interact with the system through a web-based dashboard that allows them to filter and combine data based on:

- location (group),
- sensor or node,
- type of measurement (e.g., temperature, humidity),
- time constraints.

Using these filters, users can construct queries that follow an artwork throughout its lifecycle. For example, a conservator can analyze how environmental conditions evolved while an object was stored in one room, then exhibited in another, and later placed in temporary storage. Because changing the sensor configuration (or metadata) results in a new versioned semantic descriptions, the system can correctly associate observations with their deployment context at each point in time.

A key practical benefit of the Linked Data Event Streams (LDES) publication model is that **data filtering occurs before retrieval**. Rather than fetching all historical measurements and filtering client-side, the dashboard retrieves only the relevant fragments based on semantic relations and temporal constraints. For instance, when analyzing conditions in a specific year, only the corresponding fragments are accessed, avoiding unnecessary data transfer and improving responsiveness.

This is enabled by the fact that the event stream is published as a semantically linked fragment tree. For a given query, the dashboard can follow only those fragment relations that match the selected group, sensor, and time window, and prune all other subtrees. This makes selective data access a navigation problem rather than a centralized query problem, which fits cross-institutional settings where providers

should not have to expose custom query endpoints. Instead of retrieving all fragments in Figure 2, the client filters on each level, thus only retrieving fragments along the path to the desired data. This keeps retrieval lightweight even as the overall monitoring history keeps growing.

5.2. Scenario 2: Cross-Institutional Access During Loans

A second, critical scenario concerns the monitoring of artworks during loans between museums—the situation introduced in our running example. Lending institutions typically require access to environmental data from the borrowing museum to ensure that conservation conditions meet agreed standards, while borrowing institutions must retain control over their broader monitoring infrastructure.

In MuMo, this scenario is supported through **group-based access control aligned with Solid identities** with the investigating dashboard. For a loan, the borrowing museum creates a dedicated group in the legacy dashboard and associates the relevant sensors with that group. Access to this group is then granted to specific WebID accounts belonging to the lending institution.

Because access control is enforced at the level of published data fragments, external users can authenticate using their own WebIDs and access only the data fragments corresponding to the loan group. No centralized user management or data replication is required. Once the loan period ends, sensors should be taken out of the group, so new data is not shared.

This approach enables temporary, fine-grained sharing of monitoring data across institutional boundaries while remaining manageable for museum staff and compatible with existing workflows.

5.3. Scenario 3: Integrating Multiple Decentralized Data Sources

In addition to supporting analysis within a single monitoring deployment, the advanced dashboard demonstrates the ability to **combine data from multiple independent MuMo data sources**. Each MuMo deployment publishes its monitoring data and sensor descriptions independently, yet follows the same semantic representation and event-based publication model.

In practice, this allows users to access and analyze data originating from different museum setups within a single interface. For example, a user may compare environmental conditions across multiple exhibition spaces or institutions, provided they have the appropriate access rights. Because sensor metadata and observations are published via LDES, the dashboard can discover available sensors, determine authorization, and incrementally retrieve data from multiple sources without requiring centralized aggregation.

Beyond this demonstrated functionality, the same mechanisms also enable the **conceptual integration of external data sources** that are not part of the MuMo project, such as weather station measurements. Since both sensor metadata and observations are modeled using shared semantic standards, incorporating additional event streams would not require changes to the underlying architecture. While such external integrations have not yet been deployed, they directly informed the design of the system and illustrate how MuMo supports extensibility and reuse.

5.4. Generalizing beyond monitoring systems

The scenarios above emphasize environmental monitoring data, but the broader takeaway is about interoperability between contextual datasets. Today, MuMo can derive rich analyses because it combines observations with time-aware context from the monitoring infrastructure itself (e.g., when a sensor was deployed where, and for which period). However, museums already maintain other structured context—most notably about artworks, locations, and movements—in collection management systems and repositories, such as Omeka-S or digital repository services such as D-RaaS [17].

If this object-centric information could be combined with MuMo’s sensor and observation streams in an interoperable way, MuMo would become substantially more powerful. Queries could then be expressed directly in conservatorial terms—e.g., “show me a graph of the temperature experienced by this artwork”—because the system could follow links from an object to its associated locations and time windows, and from there build queries to retrieve to the relevant measurements.

At present, the necessary datasets often exist, but they are not interoperable: a conservator can typically assemble such answers only by manually aligning exports from the collection management system or repository with monitoring data. MuMo's approach highlights that making these contextual links available in a structured, time-aware form would shift this effort from ad-hoc manual integration to repeatable, queryable reuse at the point of use, without requiring museums to replace their existing collection management infrastructure.

6. Observations and Lessons Learned

The in-use scenarios presented above highlight how specific architectural and modeling decisions shaped the practical use of MuMo in museum monitoring contexts. Rather than evaluating individual technologies in isolation, this section reflects on how these choices interacted with real-world constraints and practices.

6.1. Solid as an Institutional Boundary Mechanism

Across all scenarios, Solid played a central role in enabling data sharing without centralization. By treating Solid Pods as **institution-level data endpoints** rather than user-centric storage, MuMo aligned naturally with museum governance structures. Institutions retain control over their monitoring data while still enabling external access when required.

This was particularly evident in the loan scenario (Scenario 2), where Solid identities (WebIDs) allowed users from different institutions to authenticate without relying on a shared user database. Access could be granted and revoked by modifying group membership, supporting temporary collaborations without introducing new identity management workflows.

Lesson learned: Deploying Solid pods aligned with organizational boundaries and responsibilities enables effective data sharing without centralization.

6.2. Linked Data Event Streams for Scalable Analysis

The use of Linked Data Event Streams proved essential for handling the continuous and growing nature of monitoring data. In the analysis scenario (Scenario 1), LDES enabled the dashboard to retrieve only the relevant data subtrees based on semantic relations and temporal constraints, avoiding the need to fetch complete datasets and filter client-side.

Lesson learned: Event-based publication is well suited for long-running Digital Humanities data collection, where datasets evolve continuously.

Meanwhile, Solid access control can be enforced on the same subtree constraints. This combination of LDES and Solid allows selective cross-institutional access without requiring providers to implement bespoke query endpoints or replicate data into centralized stores.

Lesson learned: A fragmentation strategy that serves both performance needs (subtree filtering) and governance needs (selective access) makes for an effective combination of LDES and Solid in cases where selective access is more important than expressive querying.

This incremental access model also facilitated multi-source analysis (Scenario 3), where data from multiple independent MuMo deployments could be consumed and combined without requiring centralized aggregation.

6.3. Group-Based Access Control as a Deliberate Design Choice

MuMo deliberately adopts **group-level access control**, even though finer-grained authorization is technically feasible within the chosen architecture. The use of Linked Data Event Streams combined with a hierarchical fragmentation strategy makes it possible to enforce access constraints at different levels of granularity, such as individual days or, with minor adaptations to the publication pipeline, per measurement type.

In practice, these options were not activated. The legacy dashboard that remains authoritative for configuration and access management does not provide mechanisms to express such fine-grained permissions. As a result, only those authorization concepts that are meaningful and manageable within existing museum workflows were reflected in the generated access control configurations.

This decision was particularly evident in the loan scenario, where group-level access proved sufficient to grant lending institutions visibility into relevant monitoring data without exposing unrelated information. Introducing finer-grained access control would have increased technical complexity without corresponding benefits for end users. Designing for feasible access control proved more valuable than designing for maximal access control flexibility.

Lesson learned: In this applied Digital Humanities setting, the appropriate granularity of access control is determined less by technical capability than by the expressive power of existing organizational tools and practices.

6.4. Semantic Modeling as an Enabler of Integration

A key benefit of grounding MuMo in Linked Data and Semantic Web standards—RDF, SSN/SOSA, LDES, and the TREE specification—is that the resulting data is self-describing and machine-interpretable by any standards-compliant client. Semantic representations based on shared observation models enabled MuMo to separate **data production** from **data consumption**. Data is produced with two datasets: versioned sensor descriptions and observations (the actual data). These datasets are consumed together, resulting in the full picture for each observation.

This separation was critical in Scenario 1, where sensor relocations needed to be reflected correctly in longitudinal analyses, and in Scenario 3, where data from multiple sources could be combined as long as they adhered to the same semantic structures.

Lesson learned: Semantic models that inherently support independent evolution improve reuse across independently managed systems.

6.5. Integration at the Point of Use: Client-Side vs. Server-Side

Finally, the combination of Solid, LDES, semantic modeling, and client-side aggregation reflects a broader dataspace principle: **integration happens at the point of use**, not through centralized infrastructure.

An alternative approach would be server-side integration, for example through a centralized SPARQL endpoint or a federated query service that aggregates data from multiple sources before presenting it to the client. While such approaches can offer richer query expressiveness, they introduce a single point of control (and failure), require participating institutions to expose query-compatible endpoints, and shift governance to whichever party operates the central service. In museum practice—where institutions insist on controlling their own data and collaborations are often temporary—server-side integration conflicts with the requirement for institutional autonomy.

MuMo's client-side approach avoids these tensions: each institution publishes its data independently via LDES on its own Solid Pod, and consumers integrate only what they need, when they need it. The trade-off is that complex cross-source queries must be composed at the client, but for the monitoring use cases encountered in practice (time-series visualization, threshold checking, longitudinal comparison), this has proven sufficient.

Lesson learned: Point-of-use integration is a good fit for settings with independently published data under distributed governance, as it avoids the need for centralized infrastructure while remaining compatible with legacy systems and institutional autonomy, aligning with dataspace concepts.

7. Implications for the Digital Humanities Community

It is important to acknowledge that environmental sensor data is, strictly speaking, operational infrastructure data rather than cultural or humanistic content. Museum monitoring captures physical

conditions—temperature, humidity, light—for preservation management, not for humanistic interpretation or scholarly inquiry in the traditional sense. The contribution of MuMo to Digital Humanities is therefore not a new dataset for humanistic scholarship, but rather a tested architectural pattern and a set of practical lessons that are transferable to DH settings with similar structural characteristics.

Many Digital Humanities projects share the conditions that motivated MuMo: data that is collected continuously over long periods, embedded in local institutional infrastructures, and subject to distributed governance—whether sensor readings, annotations, evolving metadata, or digitized collections. In practice, such data often remains siloed within project-specific systems, limiting reuse, comparability, and shareability across institutional boundaries. Digital Humanities researchers have increasingly emphasized that the sustainability of digital resources is not only a technical issue but also an institutional and socio-technical one, shaped by funding horizons, maintenance practices, and governance constraints [18, 7].

The experience reported in this paper suggests that **dataspace-oriented approaches** are well suited to such settings. By allowing data to remain under the control of the institution that produces it while enabling integration at the point of use, dataspace-oriented approaches provide a viable alternative to centralized platforms. This pattern—decentralized publication, semantic alignment, and governance-aware access control—is not specific to sensor data and could be applied to DH collaborations involving multiple partners with heterogeneous governance structures and long-lived legacy systems.

The use of **semantic representations as infrastructural connectors**, rather than as expressive modeling exercises, further supports interoperability in constrained environments. Lightweight semantic alignment enables independently managed datasets to be combined, extended, and reinterpreted over time, without requiring uniform tooling or deep semantic expertise from end users.

Together, these observations point toward a pragmatic path for DH infrastructures that emphasizes decentralization, semantic interoperability, and governance-aware design over comprehensive but fragile integration solutions. While MuMo itself addresses an operational concern in cultural heritage, the architectural lessons are relevant wherever long-running, distributed data must be shared under institutional constraints.

8. Related work

Environmental monitoring is a well-established concern in museum practice, since temperature, humidity, and light exposure directly affect object preservation. Prior work has explored wireless and low-cost monitoring infrastructures for museums, showing the value of continuous sensing but largely focusing on local data capture and dashboard use rather than interoperable publication or cross-institutional sharing [12, 15].

A second line of related work concerns semantic sensor modeling. MuMo builds on OSLO and SSN/SOSA to represent sensors, observations, and their changing context in an interoperable way [2, 5, 10]. This is particularly relevant for museum monitoring, where sensors may be moved between rooms or reused in new configurations, while historical observations still need to remain interpretable.

For publishing continuously growing datasets, MuMo uses Linked Data Event Streams. Prior work has shown how LDES supports incremental publication and consumption of evolving linked datasets through navigable fragments [19, 16]. This fits environmental monitoring well, because observations are append-oriented and are often consumed only for a limited subset such as a specific sensor, group, or time window.

At the architectural level, MuMo is informed by dataspace research, which argues for incremental integration across independently managed sources rather than full upfront unification [8, 9]. This is closely aligned with museum practice, where monitoring data is embedded in local systems and only occasionally needs to be shared across organizations.

MuMo also relates to work on decentralized data governance through Solid, where storage, identity, and access control are decoupled from specific applications [13, 20]. Combined with Web Access Control and LDES publication in Solid containers, this supports selective sharing without requiring centralized

user management or custom query endpoints [3, 16].

Finally, prior work on authorization highlights that fine-grained access control can become complex to configure and manage in practice [11]. In that respect, MuMo's contribution is not a new access control model, but an applied architecture that aligns publication structure and authorization boundaries with existing museum workflows.

Taken together, these strands provide the foundations for MuMo, but not their combined application in an operational museum monitoring setting with legacy dashboards, continuously growing observation data, and temporary cross-institutional access needs such as loans. The contribution of this paper is therefore primarily infrastructural and applied.

9. Conclusion and Future Work

This paper reported on MuMo (Museum Monitoring), a three-year applied research project that explored the use of dataspace principles for environmental monitoring in museum practice. By integrating semantic data publication, event-based access, and decentralized governance with existing monitoring infrastructure, MuMo addressed practical challenges related to long-running data collection, cross-institutional collaboration, and controlled data sharing.

A central outcome of the project is the demonstration that dataspace-oriented architectures can be deployed incrementally alongside legacy systems (RQ1). Concretely, the in-use scenarios (Section 5) show that semantic publication via LDES, combined with Solid-based governance, supports longitudinal analysis (Scenario 1), selective cross-institutional sharing during loans (Scenario 2), and client-side integration across independent deployments (Scenario 3)—addressing RQ2. The lessons learned (Section 6) further illustrate that design choices grounded in conservatorial practice—such as group-based access control and client-side integration rather than server-side aggregation—can be both sufficient in practice and easier to sustain than more fine-grained or centralized alternatives (RQ3).

In addition, while the system supports aggregation across multiple MuMo deployments, integrating external data sources beyond the project has not yet been explored in deployed settings.

In conclusion, MuMo provides an in-use perspective on how Solid-based data governance, Linked Data Event Streams, and lightweight semantic modeling can support data management in cultural heritage institutions. While the monitored data is operational rather than humanistic in nature, the architectural patterns and lessons—decentralized publication, governance-aware access control, and semantic interoperability with legacy systems—are transferable to Digital Humanities projects that face similar challenges of long-running, distributed data under institutional governance.

Declaration on Generative AI

During the preparation of this work, the authors used Claude Sonnet to support the writing process. In particular, the tool was employed for paraphrasing and rewording text, improving writing style, and providing formatting assistance. All scientific content, including the ideas, methodology, and conclusions, was developed by the authors. The authors carefully reviewed and edited all AI-generated suggestions and take full responsibility for the final content of the manuscript.

References

- [1] Environmental monitoring, December 2025. URL https://www.conservation-wiki.com/wiki/Environmental_Monitoring. MediaWiki page. Last edited 30 December 2025.
- [2] Raf Buyle, Laurens De Vocht, Mathias Van Compernelle, Dieter De Paepe, Ruben Verborgh, Ziggy Vanlিশout, Björn De Vidts, Peter Mechant, and Erik Mannens. Oslo: Open standards for linked organizations. In *Proceedings of the international conference on electronic governance and open society: Challenges in Eurasia*, pages 126–134, 2016.

- [3] Sarven Capadisli, Tim Berners-Lee, Ruben Verborgh, and Kjetil Kjernsmo. Solid protocol, December 2022. URL <https://solidproject.org/TR/2022/protocol-20221231>.
- [4] Story Henry Capadisli Sarven and Tim Berners-Lee. eb access control (wac). Technical report, MIT CSAIL, 2024. URL <https://solidproject.org/TR/wac>.
- [5] Michael Compton, Payam Barnaghi, Luis Bermudez, Raul Garcia-Castro, Oscar Corcho, Simon Cox, John Graybeal, Manfred Hauswirth, Cory Henson, Arthur Herzog, et al. The ssn ontology of the w3c semantic sensor network incubator group. *Journal of Web Semantics*, 17:25–32, 2012.
- [6] European Commission. Linked data event streams (ldes) | semic support centre | interoperable europe portal, 2026. URL <https://interoperable-europe.ec.europa.eu/collection/semic-support-centre/linked-data-event-streams-ldes>. Accessed: 2026-01-20.
- [7] Katrina Simone Fenlon. Sustaining digital humanities collections: challenges and community-centred strategies. *International Journal of Digital Curation*, 15(1):13–13, 2020.
- [8] Michael Franklin, Alon Halevy, and David Maier. From databases to dataspace: a new abstraction for information management. *ACM Sigmod Record*, 34(4):27–33, 2005.
- [9] Alon Halevy, Michael Franklin, and David Maier. Principles of dataspace systems. In *Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 1–9, 2006.
- [10] Armin Haller, Krzysztof Janowicz, Simon J. D. Cox, Maxime Lefrançois, Kerry Taylor, Danh Le Phuoc, Joshua Lieberman, Raúl García-Castro, and Claus Stadler. Semantic sensor network ontology. W3c recommendation, W3C, 2017. URL <https://www.w3.org/TR/vocab-ssn/>.
- [11] Vincent C Hu, David Ferraiolo, Rick Kuhn, Adam Schnitzer, Kenneth Sandlin, Robert Miller, Karen Scarfone, et al. Guide to attribute based access control (abac) definition and considerations. *NIST special publication*, 800(162):1–54, 2014.
- [12] Jaime Laborda, Ana María García-Castillo, Ricardo Mercado, Andrea Peiro-Vitoria, and Angel Perles. From concept to validation of a wireless environmental sensor for the integral application of preventive conservation methodologies in low-budget museums. *Heritage Science*, 10(1):1–17, 2022.
- [13] Essam Mansour, Andrei Vlad Samba, Sandro Hawke, Maged Zereba, Sarven Capadisli, Abdurrahman Ghanem, Ashraf Abounaga, and Tim Berners-Lee. Solid: A platform for decentralized social applications based on linked data. Technical report, MIT CSAIL, 2016. URL <https://solidproject.org/TR/solid>.
- [14] Stefan Michalski. The ideal climate, risk management, the ashrae chapter, proofed fluctuations, and towards a full risk analysis model. *Experts roundtable on sustainable climate management strategies*, pages 1–19, 2007.
- [15] Laura M Rodríguez Peralta and Lina M Brito. An integrating platform for environmental monitoring in museums based on wireless sensor networks. *International Journal on Advances in Networks and Services*, 3(1/2):114–124, 2010.
- [16] Wout Slabbinck, Ruben Dedecker, Sindhu Vasireddy, Ruben Verborgh, and Pieter Colpaert. Linked data event streams in solid ldp containers. In *Managing the Evolution and Preservation of the Data Web 2022*, volume 3339, pages 28–35. CEUR, 2023.
- [17] Lefteris Tsipi, Demosthenes Vouyioukas, Georgios Loumos, Antonios Kargas, and Dimitrios Varoutas. Digital repository as a service (d-raas): Enhancing access and preservation of cultural heritage artifacts. *Heritage*, 6(10):6881–6900, 2023.
- [18] Joanna Tucker. Facing the challenge of digital sustainability as humanities researchers. *Journal of the British Academy*, 10:93–120, 2022.
- [19] Dwight Van Lancker, Pieter Colpaert, Harm Delva, Brecht Van de Vyvere, Julián Rojas Meléndez, Ruben Dedecker, Philippe Michiels, Raf Buyle, Annelies De Craene, and Ruben Verborgh. Publishing base registries as linked data event streams. In *International Conference on Web Engineering*, pages 28–36. Springer, 2021.
- [20] Verstraete, Melanie and Verbrugge, Sofie and Colle, Didier. Solid : enabler of decentralized, digital platforms ecosystems. In *ITS 31st European Conference 2022, Proceedings*, page 19, 2022. URL <https://itseurope.org/2022/>.