POSITION PAPER

Beyond FAIR: The Missing Human Sovereignty Layer in Scientific Data Infrastructure

M.E. 'Nara' Lau, Founder & CEO · Kevin W. Hartig, Co-Founder & CTO
FISE Technologies · fisetech.ai · 2026

Executive Summary

The 2016 FAIR Guiding Principles established a foundational framework for scientific data governance. A decade later, their most critical requirements including authenticated access, enforceable provenance, and human-controlled data remain unimplemented. [1]The NIH's proposed Controlled-Access Data Policy,[2] introduced in December 2025, has exposed this gap at scale: most repositories lack the infrastructure to comply. [3]This paper argues that the missing prerequisite is a human sovereignty layer, built at the protocol level, not the policy level. FISE Technologies is building that layer, with an interoperable, standards-based architecture that makes identity, access, and data governance enforceable by protocol and process design. The paper identifies three structural gaps FAIR could not close, presents the technical requirements for a sovereignty layer that closes them, and offers a set of recommendations for researchers, institutions, and policy makers.

What FAIR Got Right and What It Left Open

When the FAIR principles were published in *Scientific Data* in 2016, [4] they represented a genuine step forward. For the first time, a broad coalition of academic, industry, and government stakeholders agreed on a shared framework: data should be Findable, Accessible, Interoperable, and Reusable by both humans and machines.[5]

The principles were prescient about several things. They recognized that authentication and authorization would be necessary (Principle A1.2) and demanded detailed provenance (R1.2).[6] FAIR acknowledged that machine agents would increasingly act on behalf of humans in data ecosystems, and explicitly noted that the existing digital infrastructure was preventing researchers from extracting maximum benefit from their investments.

But FAIR also made a deliberate choice by focusing on data objects and repositories, not on the humans who generate and own the data. Authentication and provenance were listed as requirements, never implemented as standards. A decade later, that gap has widened without implemented protocols.

---

[1]Wilkinson, M.D. et al. "The FAIR Guiding Principles for scientific data management and stewardship." Scientific Data 3, 160018 (2016). https://doi.org/10.1038/sdata.2016.18

[2]National Institutes of Health. "NIH Controlled-Access Data Policy and Proposed Revisions to NIH Genomic Data Sharing Policy." Federal Register, Vol. 90, No. 241, December 18, 2025, pp. 59131–59135. FR Doc No: 2025-23246. https://www.govinfo.gov/content/pkg/FR-2025-12-18/html/2025-23246.htm

[3]Wilkinson, M.D. et al. "The FAIR Guiding Principles for scientific data management and stewardship." Scientific Data 3, 160018 (2016). https://doi.org/10.1038/sdata.2016.18

[4]Wilkinson et al. (2016): The Principles were designed to govern all digital research objects including data, code, and workflows for both human and machine actors engaged in knowledge reuse.

[5]Wilkinson et al. (2016), Principle A1.2 stipulates that protocols must support authentication and authorization where access controls are needed.

[6]Wilkinson et al. (2016), Principle R1.2 requires that data be traceable to their origin through detailed provenance records.

The Three Gaps FAIR Could Not Close

Authentication and authorization were never solved at the human layer.

FAIR principle A1.2 acknowledged authentication and authorization as necessary, but every solution built since 2016 implements access controls at the institutional or repository level, not at the level of the human data originator. When the NIH's proposed Controlled-Access Data Policy arrived in December 2025, requiring researchers to authenticate every requester and manage access by country, most repositories discovered they had no infrastructure to comply.[7] The result was not a compliance failure, it was an exposure of the missing layer.

Provenance and data lifecycle enforcement remain aspirational.

FAIR calls for detailed provenance under R1.2, but provides no mechanism to track what happens to data after it leaves a repository. There is no immutable record of downstream reuse, no enforcement of license terms, and no legal recourse for the data originator when their data is reshared outside agreed terms. The data lifecycle ends when it is uploaded; what happens after is ungoverned and unenforceable. NIH's Controlled-Access Data Policy requires protection throughout the entire data lifecycle and mandates restrictions on subsequent use, making enforceable provenance a compliance necessity, not a best practice.[8] Yet the infrastructure to meet that requirement does not exist in current repository architecture.

Human sovereignty was explicitly out of scope.

The original FAIR paper stated clearly that its principles were distinct from initiatives focused on the human scholar. FAIR was designed for machine-actionability and institutional stewardship. The human as originator, owner, and rights-holder of the data they generate was not the subject of FAIR. As AI, BCI, and neurotechnology now produce data that is intimately human, including brain signals, epigenomic profiles, and behavioral patterns, this omission has become a critical vulnerability.

The Missing Primitive

FAIR's unfinished work points to a single missing prerequisite: a human sovereignty layer built at the protocol level. Not a compliance checklist. Not a repository policy. A foundational infrastructure layer that gives every human a portable, user-owned digital identity; one that controls who accesses their data, under what terms, with legally enforceable protections, and with automated, auditable provenance from origination to end of lifecycle.

This layer must provide:

- Authenticated, human-owned identity with portable, decentralized, and user-controlled identification via a master DID (Decentralized Identifier).[9]
- Delegated access controls where the data owner sets terms once, and the infrastructure enforces them automatically.
- Smart contract-based data transactions that are legally enforceable, peer-to-peer, with terminating lifecycle endpoints.
- Immutable provenance featuring an on-chain record of every data transaction that makes unauthorized resharing discoverable and legally actionable.

---

[7] Reardon, S. "Neuroscientists challenge NIH's proposed human data access policy." The Transmitter, 2026. https://www.thetransmitter.org/data-sharing/neuroscientists-challenge-nihs-proposed-human-data-access-policy/

[8] National Institutes of Health. "NIH Controlled-Access Data Policy and Proposed Revisions to NIH Genomic Data Sharing Policy." Federal Register, Vol. 90, No. 241, December 18, 2025, pp. 59131–59135. FR Doc No: 2025-23246. https://www.govinfo.gov/content/pkg/FR-2025-12-18/html/2025-23246.htm

[9] World Wide Web Consortium (W3C). "Decentralized Identifiers (DIDs) v1.0." W3C Recommendation, July 2022. https://www.w3.org/TR/did-core/

- Interoperable, standards-based architecture compatible with W3C DIDs, Verifiable Credentials,[10] and existing FAIR infrastructure.[11]

FISE: The Infrastructure NIH's Controlled-Access Policy Requires

FISE (Foundational Identity and Sovereignty Engine) is the only framework building the infrastructure NIH's Controlled-Access Data Policy requires, at the protocol level, with an interoperable, standards-based architecture designed to work across AI agents, BCI devices, and neurotech infrastructure.[12]

Where the policy identifies the requirement, FISE provides the enforcement mechanism. Where NIH mandates authenticated access, FISE delivers user-owned, portable identity. Where NIH requires provenance, FISE delivers an immutable, legally enforceable data lifecycle record with terminating endpoints. Where NIH demands controlled licensing, FISE delivers smart contract-based peer-to-peer data agreements with automated revenue sharing at 70% to the human data originator.

FISE does not compete with FAIR, it completes it — and it does so on the terms NIH's policy now makes unavoidable.

FISE's architecture is built on open, interoperable standards including W3C Decentralized Identifiers, Verifiable Credentials, and smart contract protocols, ensuring it operates across existing repository infrastructure rather than replacing it.[13] The framework is designed to be adopted incrementally: repositories can implement the sovereignty layer without dismantling existing open access workflows.

Addressing the Objections

Will a sovereignty layer add too much friction for researchers?

This is the most common and legitimate concern. The NIH policy debate has made clear that researchers do not want to become compliance officers. The FISE framework is designed specifically to address this: access controls are set once by the data originator and enforced automatically by the protocol. The researcher does not review every access request manually, the infrastructure does. The goal is to reduce friction, not add it, by removing the institutional gatekeeping layer and replacing it with automated, user-defined permissions.

Are smart contracts legally enforceable across jurisdictions?

Smart contract enforceability is a developing area of law, and it varies by jurisdiction. FISE's approach does not rely on smart contracts as the sole enforcement mechanism. Rather, the smart contract creates an immutable, auditable record of agreed terms and data transactions that provides the evidentiary foundation for legal action in traditional courts. Legal enforceability is strengthened, not replaced, by the on-chain record. As jurisdictions continue to develop digital contract frameworks, the immutability of the record becomes a legal asset rather than a liability.

Does decentralized identity introduce new security risks?

---

[10] World Wide Web Consortium (W3C). "Verifiable Credentials Data Model v1.1." W3C Recommendation, March 2022. https://www.w3.org/TR/vc-data-model/

[11] Wilkinson et al. (2016): The Principles were designed to govern all digital research objects including data, code, and workflows for both human and machine actors engaged in knowledge reuse.

[12] Lau, M.E. and Hartig, K.W. "Beyond FAIR: The Missing Human Sovereignty Layer in Scientific Data Infrastructure." FISE Technologies, 2026. Forthcoming at fisetech.ai.

[13] International Data Spaces Association. "IDS Reference Architecture Model." Version 4.0, 2022. https://docs.internationaldataspaces.org

Any infrastructure layer introduces security considerations, and decentralized identity is no exception. The critical difference is that centralized infrastructure concentrates risk: a single breach at the repository level compromises all users. A decentralized sovereignty layer distributes that risk; each human's identity and keys are portable and individually controlled, so a breach at the platform level does not cascade to the data originator. The architecture is designed to reduce systemic risk, not eliminate all risk.

Recommendations

The following recommendations are addressed to three distinct audiences.

For researchers and the scientific community

- Engage with FISE Technologies to participate in the development of protocol-level data sovereignty standards designed specifically for neuroscience and biomedical data governance requirements.
- Advocate for protocol-level data governance requirements in responses to the NIH Controlled-Access Data Policy comment process, distinguishing between policy-layer mandates and infrastructure-layer solutions.
- Pilot sovereignty-layer infrastructure in partnership with FISE regarding existing neurodata repositories such as OpenNeuro and Neurodata Without Borders to evaluate compliance with emerging access requirements without dismantling open science workflows.

For research institutions and repositories

- Conduct infrastructure readiness assessments against the NIH's proposed controlled-access requirements, identifying specifically where authentication, authorization, and provenance capabilities are absent.
- Evaluate interoperable, standards-based sovereignty frameworks as compliance infrastructure rather than building proprietary access control systems that will not scale across repository boundaries.
- Engage with FISE Technologies for pilot integration of the human sovereignty layer into existing repository architecture.

For policy makers

- Distinguish between policy-layer governance, which establishes requirements, and protocol-layer governance, which makes those requirements technically enforceable. NIH's proposed policy identifies the right requirements; it does not specify the infrastructure to meet them.
- Support the development of open, interoperable identity and provenance infrastructure as a public good for scientific data governance, analogous to the investment in open data repositories themselves.
- Consider the human data originator, not only the institution, as a rights-holder in data governance frameworks, particularly as BCI and neurotechnology generate increasingly intimate human-origin data.

Conclusion

FAIR established the destination. A decade of progress has moved the field meaningfully toward findability, interoperability, and metadata standards. But the most consequential requirements FAIR identified including authenticated access, enforceable provenance, and human-controlled data remain structurally unsolved because they require infrastructure that surpasses FAIR's scope.

The NIH Controlled-Access Data Policy has made this gap impossible to ignore. Repositories that have operated on open access assumptions now face compliance requirements they cannot meet without a fundamental rearchitecture of how identity and access are managed.

Open, reproducible, and ethically grounded science cannot advance without first solving identity and data sovereignty at the human layer. FAIR proved that the problem is real. NIH's Controlled Data restrictions continue to expose the missing infrastructure. FISE's human sovereignty layer provides that infrastructure; making identity, access, and data governance enforceable by protocol and process design, not by mere policy.