Evidence Does Not Implement Itself: Why Science Needs Stewards at the Policy-Implementation Seam

To protect scientific infrastructure from administrative indifference, institutions must name someone responsible for carrying evidence into decisions and back again.

By Anthony Veltri

The Room Where It Wasn't

When the U.S. Department of Agriculture announced in March 2026 that it would close 57 of its 77 Forest Service research and development facilities, the scientific community responded the way scientific communities respond to existential threats: it mobilized. Researchers testified. Advocacy organizations published defenses of specific sites. Congressional offices fielded calls from constituents who understood, often viscerally, what a century of continuous forest observation was worth.

The mobilization worked, at least partially and at least for now. It also addressed the wrong problem. The community successfully defended the science after the fact, but the evidence itself was never in the room where the decision was made.

The threat to Hubbard Brook Experimental Forest in New Hampshire, one of the longest-running ecological research sites in North America, was not primarily a lack of funding, though funding was at stake. It was not primarily a communication failure, though the scientific community had to mobilize rapidly to make the case. The threat was architectural: no named person or institution was responsible for carrying the scientific value of Hubbard Brook's physical archive into the room where the closure decision was made.

Hubbard Brook's archive holds more than 60,000 barcoded samples of water, soil, plant material, and physical cores spanning seven decades of continuous collection. Nobody collecting water samples there in 1963 was thinking about PFAS contamination. But those physical samples are why scientists can now ask retrospective questions about contaminants such as PFAS, questions that were not part of the original sampling design and would otherwise be impossible to reconstruct. That retrospective capacity, the ability to ask future questions of past reality, is the archive's deepest scientific value. You cannot recreate a 1963 sample in 2026.

When the decision to close 57 facilities was made, that value had no institutional representative. Whatever internal assessment USDA may have conducted, no governance requirement existed compelling decision-makers to account for 60 years of archival option value before acting. Vicki Christiansen, the 19th chief of the Forest Service, said so publicly: the data and analysis underlying the reorganization had not been made available. The archive existed. The evidence was real. It had no route into the room.

To be clear about what failed and what did not: the scientists maintained the archive. For seven decades, they collected, catalogued, and preserved a physical record of extraordinary continuity and scientific value. The missing function was not scientific competence, commitment, or

communication. It was a formal route into the administrative decision process before the decision was made.

Hubbard Brook is the case study, not the exception. Of the 84 sites in the federal Experimental Forests, Ranges, and Watersheds network, how many have a named institutional representative whose function is to carry the scientific value of that site into the budget decisions that determine its future? Few if any appear to have such a function. The same question applies across the Long-Term Ecological Research network's more than two dozen research programs, across AmeriFlux's hundreds of monitoring locations across North America, and across the hundreds of physical archives and long-duration sampling programs that constitute the United States' scientific infrastructure for understanding environmental change. Hubbard Brook is visible because it is documented and because it has a private foundation capable of public advocacy. Most of these sites have neither. They are operating on culture rather than contract, on goodwill rather than governance, and on the assumption that their value is self-evident to the people making resource decisions. That assumption has always been fragile. It is now visibly failing.

The Seam Nobody Owns

Every major scientific institution has invested heavily in the upstream infrastructure of evidence. Data portals. Open access mandates. Knowledge commons. Persistent identifiers. Long-term monitoring networks. The investment is real and it has produced genuine advances in the availability and accessibility of scientific evidence.

And yet evidence still stalls. It stalls not because it is absent, inaccessible, or poorly communicated. It stalls because no one owns the seam between knowledge and decision.

The seam is the translation zone where evidence becomes operational: where a watershed chemistry record becomes a regulatory baseline, where a fire behavior dataset becomes a resource allocation model, where a decades-long temperature series becomes a climate adaptation plan. The seam is where science stops being a product and starts being an input to a decision someone has to make under time pressure with incomplete information and competing priorities.

Scientists are trained to produce evidence. Agencies are structured to make decisions. Nobody is formally responsible for the translation between them. This is not for want of institutional effort. Since 2019, the Foundations for Evidence-Based Policymaking Act has required CFO Act agencies to designate Chief Evaluation Officers, Chief Data Officers, and Statistical Officials, and to develop Learning Agendas and Evaluation Plans. These are genuine governance advances at the enterprise level. They do not, however, reach the site and program level of long-term research infrastructure. A Chief Evaluation Officer at USDA headquarters does not constitute a named function at Hubbard Brook with standing in the budget process that determines Hubbard Brook's future. The gap is not between the agency and the world. It is between enterprise governance and the facility-level translation function that remains unnamed.

A data management plan answers a specific question: where will the data go, how will it be documented, how will it be shared, and how will it remain accessible? That is a legitimate and

important governance requirement. But it answers a different question than the one this piece is raising. Stewardship continuity answers this: who is responsible for preserving the scientific asset's meaning across administrative decisions? That includes physical samples, custody chains, environmental controls, interpretive history, staffing, facilities, and the formal route into decisions about all of them. Federal governance systems already have categories for records, equipment, facilities, and property disposal. Long-term scientific archives do not fit neatly into any of them. A sample archive is not just a record. The freezer is not just equipment. The building is not just space. The value exists in the governed relationship among sample, method, context, environmental control, and future interpretability. Data management preserves outputs. Stewardship continuity preserves the conditions that make those outputs scientifically meaningful over time.

The few institutions that have addressed stewardship continuity directly did so deliberately, after hard experience. GRUAN, the GCOS Reference Upper-Air Network, uses mandatory site operating protocols, documentation, metadata, reporting, and site-representative requirements to protect reference-quality climate observations. ICOS, the Integrated Carbon Observation System, uses formal station labeling, evaluation, PI and staff commitments, and General Assembly approval to make participation in the network more than informal scientific goodwill. NEON operates through a cooperative agreement that assigns large-scale infrastructure management responsibilities to an identified operator.

The obvious question is whether scientific communities could develop these governance norms organically, without external mandate. Elinor Ostrom's Nobel Prize-winning work established that communities of users can often sustain shared resources through self-developed rules, without privatization or direct state control. But those arrangements depend heavily on conditions that long-term scientific archives often lack. The stakeholder community must be present: the researchers who will use a 1963 water sample to establish a PFAS baseline in 2035 cannot advocate for its preservation today because they are not yet part of any identifiable community. The governance challenge also differs in structure. At Hubbard Brook, USDA holds budget authority over a facility whose future scientific users may have no standing in the decision process. Ostrom identified recognition by external governmental authorities as a foundational condition for sustainable commons governance. That recognition is precisely what is absent. Internal community self-organization matters, but it cannot substitute for external recognition when administrative authority sits elsewhere.

These systems are not identical to Hubbard Brook or to each other, and the argument here does not require them to be. Their relevance is not that they provide a single governance model to copy. Their relevance is that each demonstrates the same governance principle: where continuity is recognized as mission-critical, mature systems make institutional responsibility explicit before failure rather than scrambling to reconstruct it after. Most long-term scientific archives in the United States have not yet made that move. They are one reorganization, one budget cycle, or one administrative decision away from the failure mode that nearly destroyed Hubbard Brook's record. They simply have not experienced the pressure yet.

Three recurring failure patterns account for most of the breakdown. Each is distinct. Each requires its own governance response. And none of them is solved by the interventions currently on offer.

Failure Pattern One: The Named Ownership Gap

The first and most fundamental failure is the absence of a named institutional owner for the seam itself.

When practitioners inside a scientific agency hear this observation, the initial response is almost always the same: but everything is owned. There are program managers. There are agency liaisons. There are communications offices and outreach coordinators and data stewards. The seam looks occupied from the inside because people are standing near it.

Proximity is not ownership. Standing near the seam and owning the seam are different things. The distinction matters because ownership implies accountability: a named person or institution that can be held responsible when evidence fails to reach a decision that needed it.

Consider what ownership of the seam would have required at Hubbard Brook. It would have required someone whose formal institutional function was to understand what the physical archive was worth to future decisions, translate that value into language the decision-making process could use, and ensure that translation reached the people making the decision before the decision was made. Not after the advocacy campaign. Not after the press coverage. Before.

That role does not exist at Hubbard Brook. It does not exist at most of the 84 sites in the federal Experimental Forests, Ranges, and Watersheds network. It does not exist as a standard function in most federal science agencies. The seam is near people, but it is not owned by anyone.

The consequence is predictable. When a reorganization reaches a research facility, the scientific value of what that facility produces enters the decision process only if someone with standing in the decision-making structure happens to know about it and chooses to raise it. That is advocacy, not governance. Advocacy is episodic, exhausting, and dependent on individual actors who may or may not be present, may or may not have access, and may or may not be believed. Governance is structural. It does not depend on the right person being in the room. It puts the right function in the room by design.

The named ownership gap is not visible to people inside the system because from inside the system the advocacy looks like ownership. Scientists speak up for their facilities. Program officers defend their portfolios. Communications teams issue statements. The machinery of institutional self-defense is active. What is missing is the standing role: the named function with formal authority to translate scientific value into the terms the decision process uses, before the decision is made.

The gap is not between who maintains the archive and who should maintain it. The scientists at Hubbard Brook have maintained that archive with extraordinary care for seven decades. The gap is between who maintains the archive and who was in the room when the budget was zeroed. Those are different rooms. They require different roles. Institutions have invested heavily in staffing the first room. The second room has a standing vacancy.

Failure Pattern Two: The Manual Reporting Tax

The second failure pattern is visible to everyone who has worked at the boundary between scientific and decision-making institutions. It is so familiar it is often mistaken for a natural feature of the landscape rather than a correctable governance failure.

Every budget cycle, every reorganization, every new administration produces the same demand: explain again, from the beginning, why this work matters. Scientists and program staff spend weeks or months reassembling evidence that has been assembled before, translating findings that have been translated before, making cases that have been made before. The knowledge exists. The evidence is real. The institutional memory of why it matters does not persist between decision cycles.

This is the manual reporting tax: the recurring cost of starting from zero because no standing function maintains the translation between what science produces and what decision-makers need.

Every adult who has navigated a complex medical history through a fragmented healthcare system has paid a version of this tax. A new provider, or a familiar provider who has not reviewed the record, requires the patient to reconstruct context that should already be in the room. The appointment is 15 minutes. The intake takes 10. What remains for actual diagnosis and treatment is not the appointment that was needed. It is the residue of an administrative failure.

In 2018, I was deployed to support Hurricane Florence response operations as a federal employee with the U.S. Department of Agriculture. During that deployment I contracted what was eventually diagnosed as West Nile infection, presenting initially with symptoms consistent with Guillain-Barré syndrome. The months that followed involved repeated clinical interactions that began the same way: establish the baseline, explain the history, account for examination findings that looked unremarkable on paper but represented a significant departure from where I had been before. Every new provider required the same reconstruction of context. Every appointment began where every other appointment had begun, not where the last one had ended.

The physician who eventually helped restore my muscular function was different in one specific way. He had read the case. The appointment began where the last one ended. That single change, the standing transfer of context rather than its reconstruction from scratch, changed everything about what the time together could produce.

Scale this to the budget cycle of a federal scientific program. Every new administration, every reorganization, every appropriations hearing requires the same reconstruction of context that should already be in the room. The scientists who spent six weeks preparing testimony for a budget hearing could have spent those six weeks on the science the testimony was defending. The program officers who rebuilt the case for a monitoring network from scratch could have maintained the case continuously at a fraction of the cost. The standing transfer of context is necessary but not sufficient on its own; it must be maintained by a named function with genuine standing, or it collapses back into the same episodic cycle the moment institutional attention shifts.

The tax exists because no institution has assigned standing ownership of the translation function. When evidence crosses the seam only under crisis conditions, every crossing requires the full infrastructure of advocacy: the political mobilization, the media strategy, the congressional outreach, the expert testimony, the rapid assembly of a case that should have been maintained continuously. Each crisis is treated as unique. Each response is built from scratch. In a period of rapid institutional change, the tax is being levied continuously on institutions whose capacity to pay it has been reduced by the very forces demanding payment.

Failure Pattern Three: The Orphaned Platform

The third failure pattern is subtler and in some ways the most consequential, because it operates silently in institutions that believe they have already solved the first two problems.

Long-term monitoring systems, physical archives, and scientific data platforms require more than initial investment and ongoing maintenance. They require a standing function that carries what they produce into operational decisions and carries the experience of those decisions back into the knowledge system. Without that function, the platform exists but does not connect. It produces evidence that accumulates without becoming operational. It is maintained without being used. It is funded without being governed.

This is the orphaned platform: a scientific asset launched but not governed at the translation layer, maintained but not connected to real decision authority, operated but not stewarded at the seam.

Scientists are often reluctant to occupy the translation function and the reluctance is understandable. The professional identity of a research scientist is built around producing evidence, not deploying it. When scientists are asked to engage with the policy process, the response is frequently some version of: I just do the science, I am not a policy person. This is not evasion. It reflects a genuine and well-grounded understanding of where scientific credibility comes from. A scientist who becomes an advocate risks the independence that makes their work credible. The science-policy boundary is not arbitrary. It protects something real.

But the orphaned platform failure is not caused by scientists refusing to do policy. It is caused by institutions failing to staff the translation function as a distinct role, separate from both the scientific work and the policy work, responsible for the seam between them.

The specific failure at Hubbard Brook illustrates this precisely. The scientists maintained the archive. The institution was funded. The evidence was real. No one was assigned the function of carrying the archive's value into the budget decision. The archive was not orphaned in the sense of being abandoned. It was orphaned in the sense that its connection to the decision process depended entirely on whether the right advocate happened to be in the right room at the right time.

A further dimension of this failure deserves naming. Even when evidence does cross the seam and reaches a decision, the experience of that decision rarely returns to the knowledge system. What did the decision-makers actually need? What questions did the evidence fail to answer? What did implementation reveal about assumptions that turned out to be wrong? That return path

is as important as the forward path from evidence to decision. Without it, the knowledge system stops learning what decisions need from it. It becomes increasingly sophisticated at producing evidence for questions it has already answered and increasingly blind to the questions that are actually being asked.

A Diagnostic Framework

These three failure patterns are distinct, but they share a common structure: each represents a failure at a specific layer of what a complete stewardship system requires. Understanding the layers matters because the most common institutional mistake is applying the right intervention at the wrong layer and counting it as a solution. A complete stewardship system for a long-term scientific asset requires at least five distinct functions, each with its own responsible party and its own failure mode.

The first is operational stewardship: keeping the physical or technical infrastructure running, executing collection protocols, and maintaining the institutional knowledge required to interpret what the instrument produces under non-standard conditions.

The second is custodial stewardship: maintaining the integrity and continuity of what the system produces, including the methodological chain that makes archived materials usable by people who were not present when they were created. A custodial obligation begins at first collection and does not end when the producing operation concludes. Operations can be a project. Custody cannot.

The third is internal governance design: writing stewardship obligations into the operating structure of the institution itself, through founding documents, cooperative agreements, and operating contracts. This is the layer that makes stewardship a structural requirement rather than a cultural practice. When the founding generation leaves and the culture changes, a contractual obligation survives. A cultural practice does not.

The fourth is external governance protection: the legislative, regulatory, and political instruments that protect the institution from adverse administrative action. Appropriations language. Statutory continuity requirements. Mandated transition planning before facility reduction.

The fifth, and the one most consistently absent, is translational stewardship: the named function responsible for carrying evidence across the policy-implementation seam into decisions and back again. Someone who understands what the scientific asset produces, can translate that value into the terms decision-makers use, and is formally responsible for ensuring that translation occurs before decisions are made rather than after they are challenged. This differs from what scholars call 'boundary spanning,' 'knowledge brokering,' or the 'honest broker' role in one critical respect: it requires not just the function but formalized authority, a continuity obligation that survives personnel and political transitions, and bidirectional accountability for the return path from decisions back into the knowledge system.

No layer substitutes for another. This is the central diagnostic claim. Legislative protection at the fourth layer does not solve the orphaned platform problem at the fifth. A named evidence steward at the fifth layer cannot compensate for a broken archive at the second. Each layer is

necessary. None is sufficient on its own. The most dangerous failure pattern crosses layers silently: an institution can have strong operational stewardship, intact custodial records, contractual governance obligations, and legislative protection, and still fail at the fifth layer, producing evidence that accumulates without ever reaching the decisions it was generated to inform.

The Governance Instrument

The three failure patterns share a common direction of remedy, though the specific instrument depends on which layer has failed.

For the named ownership gap, the instrument is designation: a named role with formal authority and accountability for the translation function, distinct from both the scientific work and the communications work, responsible specifically for the seam.

For the manual reporting tax, the instrument is standing maintenance: a continuously updated record of what the scientific asset produces in terms relevant to decision-makers, maintained between decision cycles rather than rebuilt from scratch at each one. This is not a communications function. It is a governance function.

For the orphaned platform, the instrument is loop closure: a formal requirement that the return path from decisions back into the knowledge system is maintained as a structural obligation equivalent to the forward path from evidence to decision.

In each case, the intervention requires naming. Not a committee. Not a working group. Not a responsibility distributed across existing functions. A named person or institution with a specific accountability, a defined scope, and a governance instrument that survives personnel turnover, budget cycles, and changes in political environment.

The durability of the governance instrument matters as much as its existence. A public commitment to maintaining a research facility is not the same as a contractual obligation to do so. A press release cannot authorize a facility manager to keep the lights on. A statutory requirement to produce a continuity plan before any facility reduction creates friction that weaker instruments cannot provide. An appropriations proviso that conditions facility reduction on Congressional approval removes administrative discretion entirely. The difference between these instruments is structural, not rhetorical.

The Role, Not the Burden: What the Steward Actually Does

The governance response described here will produce a predictable set of objections. Scientists will ask whether they are being asked to become advocates. Administrators will ask whether a new role means more overhead. Legislators will ask whether this requires new authority. Funders will ask whether it requires new budget lines.

The honest answer to each objection is the same: the named evidence steward exists specifically so that everyone else does not have to change what they do.

Before that argument can land, a prior problem must be addressed. In most institutions, when someone finally recognizes that the translational stewardship function needs to exist, the response is to assign it to whoever has capacity. It becomes an "other duties as assigned" addition to a role that was already fully defined. Someone receives the label without receiving the training, the standing, the authority, or the clear definition of what success looks like. The function is absorbed into a communications role, a program management role, or a data management role where its distinctive requirements are subordinated to the priorities of the host function. A filing cabinet drawer is labeled with the right name. The function it was supposed to contain is not there.

This is not a named steward. It is a named-steward-shaped hole. The distinction between naming a role and defining it with genuine authority and accountability is the difference between governance and the appearance of governance. Institutions that create the appearance will continue to experience the failures the role was meant to prevent, and will have the additional disadvantage of believing they have already addressed the problem.

The reflex to form a committee rather than designate a person compounds this failure. A committee can produce a report. It cannot be held responsible for whether that report reaches a decision. When accountability is distributed across a group, the individual incentive to ensure the outcome is achieved drops to near zero. Diffused accountability is not reduced accountability. It is, in practice, no accountability at all.

The institutional logic that produces these failures is not irrational on its own terms. Organizations optimize for what they can measure. Travel policy violations are legible: the receipt above the per diem rate, the upgrade that required approval, the ground transportation choice that could have been cheaper. These costs are visible, countable, and auditable. They generate policies, compliance requirements, and approval workflows precisely because their deviation is easy to detect and their value easy to calculate. The scientific value of a 60-year unbroken archive of watershed chemistry is none of those things. It cannot be entered into a spreadsheet at the moment of a budget decision. Its option value, the ability to answer questions nobody has yet thought to ask, does not appear on a balance sheet. So it does not generate policies. It generates goodwill, institutional culture, and the assumption that someone is taking care of it.

There is a particular irony in this failure when it occurs inside an agency whose statutory mission is stewardship. The U.S. Forest Service was founded on the doctrine of the wise use of natural resources across generations. It stewards 193 million acres of public land against the pressures of the immediate. It maintains forests so that future generations can use them. That doctrine, applied to the scientific record of what those forests actually do, would produce exactly the governance architecture this piece is arguing for: named responsibility, long-term commitment, and the explicit subordination of short-term administrative convenience to multigenerational institutional value. The agency knows how to steward land. It has not applied that knowledge to stewarding the evidence the land produces. The junk drawer is not a failure of values. It is a failure of institutional scope.

A genuine translational stewardship role requires definition on four dimensions before it can function. It requires a named person with a clear accountability, not a shared function or a committee. It requires standing in the budget and planning process, not merely access to it. It requires a scope that is specific enough to be auditable: what evidence, from what system, reaching what decisions, documented in what form. And it requires a governance instrument that makes the function survive beyond the individual who first holds it.

Each actor faces both a disincentive and a structural penalty for going first. The disincentive is rational: a scientist who advocates for a named steward without institutional backing risks being read as protecting their own program; an administrator who creates the role without external mandate absorbs cost that peer institutions have not accepted; a funder who requires it alone may simply redirect applicants toward funders who do not. But the structural penalty is more consequential than the disincentive.

Even if one actor overcomes the disincentive and introduces the requirement unilaterally, the requirement remains fragile. A standard that is not broadly adopted is not a standard. It is a local rule, enforced by one person, overridden by every other actor who did not adopt it, and destined to be absorbed into the "other duties as assigned" drawer the moment its champion moves on.

Military environments offer a useful analogy, and anyone who has worked in or adjacent to a military environment has observed the contrast directly. There is no sign on the sergeant major's lawn. There is no written policy that reaches civilian contractors who are technically outside the chain of command entirely. And yet no one walks on the lawn. The standard shapes behavior even in the absence of formal authority over the actor, because it is sufficiently embedded in the environment that everyone who moves through it absorbs and reinforces it regardless of their official relationship to its origin. The sergeant major does not need to be present for the standard to hold.

A coordinated external requirement for translational stewardship produces exactly this effect. A grant condition applied across the research funding enterprise makes the standard legible to every actor who moves through that environment: scientists, administrators, program officers, review panels, and institutions whose formal relationship to any individual funder is indirect at best. Non-compliance becomes visible and costly rather than invisible and costless. The standard becomes self-reinforcing precisely because it was not introduced by one actor acting alone.

With that structural reality established, consider what each actor actually does differently when a genuine named steward is in place.

Scientists do one thing differently: they share context with a trusted translator. They do not subordinate their scientific judgment to the steward. They do not modify their research to serve the steward's needs. They do not testify, brief, or advocate in roles their professional identity was not built for. They make themselves accessible enough for the steward to accurately understand what the work produces and why it matters. The steward does not change the science. The steward carries what the science produces into rooms the scientists should not have to enter.

Agency administrators do one thing differently: they create a standing seat in the budget and planning process for the translational stewardship function. Not a voice that can be invited and uninvited. A seat. The distinction matters because presence without standing produces the advocacy dynamic the governance instrument is designed to replace.

Funders do one thing differently: they add a named translational stewardship accountability to grant conditions, in the same way they already require data management plans and data sharing commitments. Name the person or institution responsible for carrying what the research produces into operational decisions. Document how that function will be maintained across the life of the award. Enforcement requires integration into mid-award review, not a filed template. The question is not whether a steward has been named but whether that person holds a standing seat in the institution's actual planning and budget process.

Legislators do one thing differently: they write the function into the operating conditions for facilities that receive public funding, in the same way they write data integrity and environmental control requirements.

The steward absorbs what everyone else currently does inefficiently and episodically: the continuous maintenance of the translation between scientific value and decision-maker vocabulary, the standing presence in budget and planning processes, the current records of what the scientific asset produces in terms the decision process uses, the return path from decisions back into the knowledge system, and the institutional memory across personnel and political transitions.

The profile required is unusual and worth naming precisely. A genuine evidence steward needs enough scientific literacy to accurately represent what the research produces. Enough policy fluency to translate that representation into the terms decision-makers use. Enough administrative experience to understand how resource decisions are made. Enough institutional independence to represent the evidence even when it is inconvenient for the institution. They are not a scientist, a communications officer, a lobbyist, or a program manager. But they need the working competence of all four simultaneously, because the seam they occupy requires all four kinds of knowledge to hold.

This profile is currently rare because institutions do not create the role and therefore do not develop people toward it. Scientists who naturally move toward the translation function are often absorbed into communications or policy positions where the translation work is subordinated to the institutional priorities of those functions. The governance gap is self-reinforcing: the absence of the role means the people who could fill it do not develop in that direction, which means when the role is finally created, the talent pipeline for it is thin. Naming the role is necessary but not sufficient. Institutions must value it visibly enough that capable people develop toward it.

Who Must Act

Three actors have the standing and the mechanisms to close the governance gap at the translational layer.

Federal agencies can designate evidence stewards for major science-to-policy initiatives through policy directives and organizational design. The function need not be created from scratch. It can be defined and formalized within existing structures, provided it carries genuine authority and genuine accountability rather than being absorbed into communications or project management roles where its distinctive function will be diluted.

Funders can require translational stewardship as a condition of support. The pattern already exists in other domains: clinical research funders require data sharing plans; environmental compliance frameworks require reporting to decision-making bodies. Extending that logic to long-term ecological research infrastructure would require research institutions to name who is responsible for carrying what their work produces into operational decisions and back.

Universities and research institutions can adjust reward structures to value boundary-spanning, maintenance, and translation work alongside novelty and publication. The current incentive structure routes the best scientific talent toward discovery and away from stewardship. That routing is rational given the current incentives. It is not inevitable.

None of these changes require legislation, though legislation would help at the fourth stewardship layer. They require institutional decisions about what functions matter, who is responsible for them, and what counts as success.

If one actor must move first, it is the funder. A grant condition requiring named translational stewardship, applied at the point of award, seeds the role at scale across institutions without requiring legislative action, organizational restructuring, or a champion inside any individual agency. Once funders require it, agencies must interface with it. Once agencies must interface with it, they develop internal capacity to match. The sequence matters. The funder mandate is where it starts.

This is not an overreach of the funder’s role: it is a standard mechanism of federal award governance. While Congress appropriates the funds, science agencies write the cooperative agreements, grant conditions, and operating requirements that shape how those funds are executed. Funders already mandate complex requirements for data management and cybersecurity, and they possess the leverage to enforce them. When the National Ecological Observatory Network’s original management structure broke down amid major cost, leadership, and oversight problems, NSF did not simply observe the failure. It replaced NEON Inc. with Battelle as the management entity and later continued Battelle’s role through a major cooperative agreement. Adding a translational stewardship requirement is not a novel administrative burden. It is an update to existing award-governance machinery so that long-term scientific assets have a responsible continuity function before the next transition or budget shock arrives.

A question arises about funder reach into intramural facilities. Hubbard Brook is a USDA site; an NSF grant condition does not directly bind USDA's internal budget process. The connection runs through a structural dependency that is often invisible until it fails. Most federal long-term research sites operate as layered systems: intramural agency funding provides baseline infrastructure while extramural grants fund the university partnerships and continental-scale data products that depend on those sites as anchor points. When funders require translational stewardship as a condition of the grants that university researchers use at a federal facility, they create a governance requirement the intramural facility must interface with in order to remain a viable partner for extramural science. The agency does not have to comply directly. It finds, over time, that it cannot maintain the research partnerships that justify the facility's existence without developing internal capacity to match.

The Archive and the Decision

Hubbard Brook's physical archive appears to have avoided the immediate closure threat, at least for now. The political mobilization worked. The scientific community made its case. The senator intervened. The secretary of agriculture gave a public commitment to reconsider.

That commitment is a reputational instrument. It creates political friction. It requires no internal operational change. A facility manager cannot use it to authorize keeping the environmental controls running. The next reorganization, the next administration, the next budget cycle will face the same structural conditions that produced the 2026 threat, because the governance gap that made the threat possible has not been closed.

The archive holds decades of continuous environmental observation. It holds samples that future scientists will use to answer questions nobody has thought to ask yet. It holds the physical reality of a watershed across seven decades of industrial pollution, acid rain, climate change, and PFAS contamination. That is not a sentimental description. It is an operational one. The archive's value is not what it tells us today. It is the fact that it preserves reality in a form future scientists can still interrogate.

That value had no route into the room where the decision was made.

Building that route is not a scientific problem. It is a governance problem. The science community already knows what the archive is worth. The governance architecture for ensuring that knowledge reaches decisions does not yet exist.

The healthcare system analogy holds at the systemic level. A system that relies on patients to reconstruct their own histories at each clinical encounter is not just inefficient. It is structurally incapable of delivering the care those histories would make possible. The patients who receive care from providers who have read the record are not luckier than the others. They are better governed.

The scientific enterprise already knows how to produce evidence. The governance has not caught up to the obligation of delivering it.

Anthony Veltri (anthony@anthonyveltri.com) is an independent practitioner and former physical scientist and senior policy analyst with the USDA Forest Service Washington Office, where he worked on enterprise architecture and governance in federal programs including those supporting scientific research. His work on the governance gap in long-term ecological archives appears in Eos (American Geophysical Union, 2026). A companion manuscript is under peer review at the Bulletin of the American Meteorological Society.

Submission History

A version of this argument was first pitched to *Issues in Science and Technology* because of its direct connection to the Dosemagen and Ottinger essay published in their Spring 2026 issue. The editors suggested the Forum section as a more immediate vehicle for engaging the Dosemagen and Ottinger readership directly, which I accepted. A shorter Forum response is in production there and will appear in the print edition and on the Issues website in late June 2026. The full

argument is published here because a complete version deserves a permanent home on its own timeline.

Sources and Recommended Reading

Shannon Dosemagen and Gwen Ottinger, "Constructing a New Knowledge Infrastructure," Issues in Science and Technology, Vol. XLII, No. 3, Spring 2026. https://issues.org/environmental-knowledge-commons-dosemagen-ottinger/

Anthony Veltri, "The Governance Gap Threatening Long-Term Ecological Archives," Eos 107 (2026), https://doi.org/10.1029/2026EO260172. https://eos.org/opinions/the-governance-gap-threatening-long-term-ecological-archives

WBUR On Point, "Is Trump taking a chainsaw to the Forest Service?" May 18, 2026. [Source for Vicki Christiansen's public statement that the data and analysis underlying the Forest Service reorganization had not been made available.] https://www.wbur.org/onpoint/2026/05/18/is-trump-taking-a-chainsaw-to-the-forest-service

Anthony Veltri, "The Governance Failure at Hubbard Brook: Why New Hampshire Must Demand a Transition Plan," New Hampshire Bulletin, April 8, 2026. https://newhampshirebulletin.com/2026/04/08/the-governance-failure-at-hubbard-brook-why-new-hampshire-must-demand-a-transition-plan/

World Meteorological Organization / Global Climate Observing System, The GCOS Reference Upper-Air Network (GRUAN): Guide, GCOS-171. Geneva: WMO, 2013. [Governance structure specifying named-stewardship obligations at upper-air reference observation sites.] https://www.gruan.org/gruan/editor/documents/gcos/gcos-171.pdf

Mitch Ambrose, "In Wake of NEON Turmoil, NSF Rethinking Its Approach to Large Facility Management," FYI: Science Policy News (American Institute of Physics), September 1, 2016. [Open access policy analysis detailing the National Science Foundation's termination of the original NEON Inc. cooperative agreement and the transfer of observatory management to Battelle following structural leadership failures. NSF's own root-cause analysis cited the need to establish "activist management credibility."] https://www.aip.org/fyi/2016/wake-neon-turmoil-nsf-rethinking-its-approach-large-facility-management

Janne Heiskanen, C. Brümmer, N. Buchmann, et al., "The Integrated Carbon Observation System in Europe," Bulletin of the American Meteorological Society 103, no. 3 (2022). [Details the European Research Infrastructure Consortium governance framework and the site-level institutional commitments required to maintain long-term stewardship across ICOS stations.] https://doi.org/10.1175/BAMS-D-19-0364.1

Paul Voosen, "Unique Canadian ice core collection suffers catastrophic meltdown," Science, April 2017. [Foundational case study on the permanent destruction of scientific capital when physical archive infrastructure fails without a formalized continuity and stewardship plan.] https://www.science.org/content/article/unique-canadian-ice-core-collection-suffers-catastrophic-meltdown

Elinor Ostrom, "*Governing the Commons: The Evolution of Institutions for Collective Action*." Cambridge University Press, 1990. [Foundational work establishing that communities of users can sustain

shared resources indefinitely through self-developed governance rules, without privatization or state control. Winner of the 2009 Nobel Prize in Economics.]

Anthony Veltri, “Federation Cannot Replace Stewardship: Governing the Calibration Chain in Continental-Scale Observation Networks,” Bulletin of the American Meteorological Society, manuscript (under peer review, 2026).

Anthony Veltri, “Interface Stewardship Doctrine,” https://anthonyveltri.com/guide/doctrine-25-the-five-stewardship-layers-a-diagnostic-taxonomy/ (2026). [Five-layer stewardship taxonomy referenced in text.]