What Is Missing Before Resources Move
The Missing Public-Interest Infrastructure for Issue-Governed Philanthropy

Thesis

Over the past two decades, philanthropy has built a remarkable collection of tools, platforms, databases, frameworks, and operating systems. Funders can search nonprofit profiles, explore grant histories, review financial information, compare organizations, manage application workflows, track portfolios, evaluate outcomes, and monitor grants with a level of sophistication that would have been difficult to imagine a generation ago.

Yet a curious problem remains.

Despite all of this infrastructure, it is still surprisingly difficult to answer some of the most important questions that arise before resources are allocated.

- Who is responding to an issue, and who is not?
- Where is funding already flowing, and where are the gaps?
- What are communities experiencing that institutions may not see?
- Which barriers matter most?
- What roles are missing?
- What does the issue itself require before specific organizations are compared?

These questions sit at the center of many philanthropic decisions, especially when funders are trying to address complex social issues, persistent disparities, place-based challenges, or long-standing conditions that have resisted change despite decades of investment. Yet the systems most commonly used in philanthropy were rarely built to answer them directly.

This paper argues that the problem is not a lack of information. Philanthropy possesses enormous quantities of information. Nor is the problem a lack of tools. The sector has invested heavily in building systems that support discovery, due diligence, grantmaking operations, reporting, evaluation, and learning.

The deeper challenge is that these systems were built around different purposes and different units of analysis. Some are organized around organizations. Others are organized around grants, places, outcomes, donors, applications, portfolios, or reporting requirements. Each contributes something valuable. Together they form an increasingly sophisticated ecosystem.

What they do not yet consistently produce is a governed issue-level view before resources are allocated.

The distinction is subtle but important.

A funder may be able to identify strong organizations. A donor may be able to compare charities. A grants manager may be able to process applications efficiently. A researcher may be able to locate evidence about a specific intervention. A community foundation may be able to see historical funding patterns.

Yet none of those capabilities automatically answers a different question:

What is happening within the issue landscape itself?
That question requires bringing together multiple forms of knowledge that are typically stored in different places, governed by different actors, and interpreted through different lenses. It requires understanding conditions of need, existing response systems, capital flows, community-level signals, contextual factors, and the roles organizations play within a larger ecosystem. It also requires determining who has authority to define, interpret, validate, and challenge what is seen.

The first paper in this series argued that philanthropic decision-making often relies on organization-level processes to answer issue-level questions. It suggested that this creates a level-and-unit-of-analysis mismatch. Organization-level review remains essential for assessing capacity, stewardship, governance, compliance, and readiness, but it cannot, by itself, provide a complete understanding of an issue landscape. The strongest evidence did not support abandoning organization-level review. Instead, it supported a different conclusion: issue-level questions require issue-level preparation.

That conclusion raises a practical question.

If issue-governed decision preparation is necessary, what would need to exist for it to become possible at meaningful scale?

This paper explores that question.

Rather than beginning with a proposed solution, it begins by examining the infrastructure philanthropy already possesses. Doing so is important because claims of novelty are often exaggerated. The sector is not operating in an information vacuum, nor is it waiting for someone to invent the first useful decision-support tool. Many capable systems already exist. They perform important functions, and any serious proposal for new infrastructure must begin by acknowledging their contribution.
Only after understanding what already exists can we meaningfully ask what remains missing.

Table of Contents

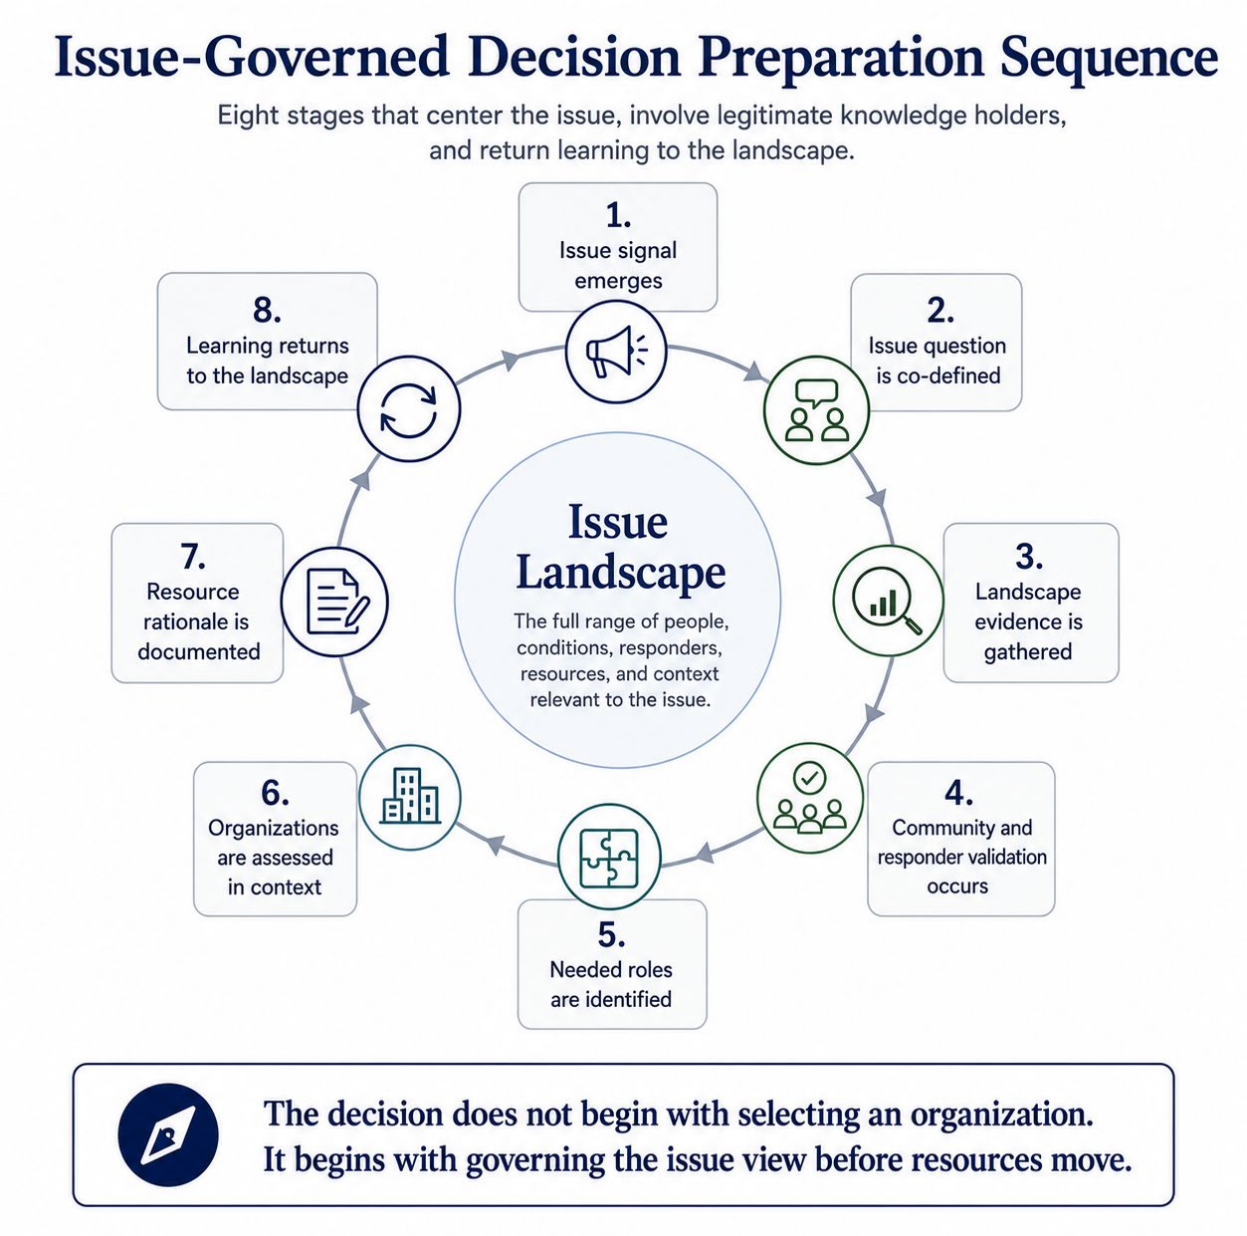

1. The Infrastructure Philanthropy Already Has

One of the most striking findings from the research supporting this paper is that philanthropy is not suffering from a lack of infrastructure. In fact, the opposite may be closer to the truth. Over time, the sector has built a substantial collection of specialized

systems designed to solve specific problems that arise before, during, and after funding decisions are made.

This is worth emphasizing because discussions about innovation in philanthropy often begin with the assumption that the sector lacks information, lacks technology, or lacks decision-support capacity. The evidence suggests something more complicated. Philanthropy possesses a growing collection of sophisticated tools. What is less clear is whether those tools can be brought together in a way that supports issue-level understanding before resources move.

The current landscape resembles a collection of highly capable instruments, each designed for a particular purpose. Some systems help users discover organizations. Others help funders process applications and manage grant portfolios. Some focus on outcomes and evidence. Others focus on geographic conditions, demographic trends, community indicators, or stakeholder engagement. Each system makes a different part of the landscape visible.

Candid provides perhaps the clearest example of infrastructure organized around organizations, funders, and grants. Its databases allow users to explore nonprofit profiles, funding histories, organizational characteristics, populations served, and grant activity across a large portion of the sector. For many users, Candid functions as the primary transparency layer of philanthropy because it helps answer foundational questions about who exists, what they do, who funds them, and how organizations compare to their peers.

Other systems approach the landscape from entirely different directions. PolicyMap organizes thousands of indicators around geography and place. Rather than beginning with organizations, it begins with conditions. Users can explore demographic patterns, economic conditions, housing characteristics, health indicators, opportunity structures, and a wide range of place-based measures. The resulting view is valuable precisely because it reveals dimensions of an issue that may never appear in an organizational application or annual report.

Impact Genome starts from yet another perspective. Its focus is not primarily on organizations or places, but on outcomes, interventions, beneficiaries, evidence quality, and comparability across programs. The question it attempts to answer is different: not who exists or where conditions are concentrated, but what kinds of interventions appear connected to particular outcomes and what evidence supports those claims.

The same pattern appears throughout the philanthropic technology landscape. Workflow systems such as Fluxx, SmartSimple, Foundant, and Submittable help grantmakers manage applications, reviews, budgets, approvals, reporting, and compliance. GiveWell conducts deep investigations into selected funding opportunities. Giving Compass helps

donors explore causes and organizations. Community listening platforms help capture stakeholder feedback. Public-sector data systems help track conditions and trends.

Viewed individually, these systems are often impressive.

Viewed together, they reveal something more interesting.

They were built around different centers of gravity.

2. What Is Still Missing Before Resources Move

If the philanthropic sector already possesses nonprofit databases, grantmaking platforms, outcome frameworks, place-based indicators, donor discovery tools, reporting systems, and growing collections of public data, then an obvious question follows.
What exactly is missing?

The answer is not immediately obvious because many of the individual components needed for informed decision-making already exist. A funder can search for organizations. A donor can compare charities. A grants manager can review applications. A researcher can locate evidence about interventions. A community foundation can examine demographic indicators and historical funding patterns. Viewed separately, each of these capabilities appears useful and often sophisticated.

The challenge becomes visible only when a decision-maker attempts to understand an issue before resources are allocated.

At that moment, the limits of existing infrastructure begin to appear.

A funder seeking to understand childhood hunger in a specific region may be able to identify nonprofits working in food access. The same funder may be able to locate census data describing poverty rates, food insecurity indicators, or demographic trends. Historical grant records may reveal which organizations have received funding in the past. Outcome frameworks may provide evidence about intervention models that have worked elsewhere. Community engagement efforts may generate stories, observations, concerns, and lived experience from affected residents.

Yet none of these sources, by themselves, answers a larger question:

What does the issue landscape currently look like?

The answer remains difficult because the relevant information is typically organized around different units of analysis.

- Some systems are organized around organizations.
- Some are organized around grants.
- Some are organized around geography.

- Some are organized around outcomes.
- Some are organized around applications.
- Some are organized around donor interests.
- Some are organized around compliance and reporting requirements.
- Each system answers an important question.
- Few answer all of them together.

This is not a criticism of existing platforms. In many cases, those systems are doing exactly what they were designed to do. Candid was built to increase visibility into organizations, funders, and grants. PolicyMap was built to make place-based conditions easier to understand. Impact Genome was built to standardize outcomes and intervention evidence. Grantmaking systems were built to help funders manage the grant lifecycle efficiently. The problem is not that these systems are failing. The problem is that they were built to solve different problems.

When viewed collectively, something important becomes visible.

The sector has become increasingly effective at understanding individual parts of an issue while remaining less effective at understanding the issue as a whole.

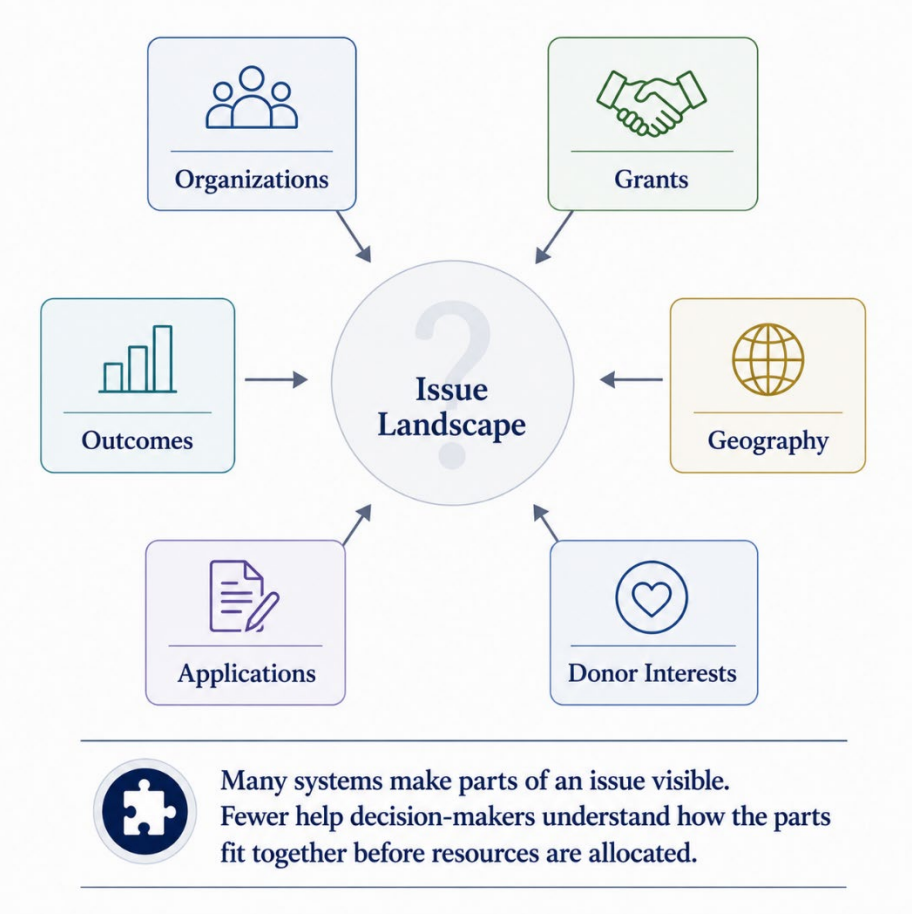

This distinction may seem subtle at first, but it becomes significant when decisions become more complex.

Choosing between two applicants is different from understanding whether a response ecosystem contains critical gaps.

Reviewing a proposal is different from understanding whether resources are already concentrated in one part of the issue landscape while other needs remain underserved. Evaluating organizational performance is different from understanding whether communities experience the broader response system as accessible, trusted, coordinated, and effective.

Examining a budget is different from understanding how capital moves across an issue over time.

The further a decision moves away from a single organization and toward a broader issue, ecosystem, population, or place, the more difficult it becomes to assemble a coherent view from existing systems.

This is where the idea of a missing layer begins to emerge.

The research conducted for this paper repeatedly identified six categories of information that decision-makers need when attempting to understand an issue before allocating resources. Interestingly, none of these categories is inherently new. All of them exist somewhere within the current philanthropic ecosystem. What is unusual is finding them assembled, governed, interpreted, and maintained together within a single decision-preparation environment.

The first category is **conditions of need**.
Most funders would agree that understanding need is foundational. Yet conditions of need are often scattered across public datasets, academic studies, local assessments, community reports, administrative records, and organizational narratives. Different sources may describe different parts of the same condition, using different definitions, geographies, timeframes, and assumptions. Understanding need requires more than locating data. It requires assembling a coherent view of what is happening, for whom, where, and with what degree of confidence.

The second category is **institutional response**.
Understanding an issue requires knowing who is already responding. This sounds straightforward, but it is often surprisingly difficult. Formal nonprofits represent only part of a response landscape. Public agencies, tribal governments, mutual aid networks, schools, clinics, faith communities, advocacy organizations, community leaders, informal support systems, and collaborative initiatives may all play important roles. Existing systems rarely

provide a comprehensive view of response ecosystems because they were not designed to map relationships, roles, and coverage across an issue.

The third category is **capital flow**.
Many philanthropic systems can show individual grants. Far fewer help decision-makers understand how resources move across an issue. Questions about concentration, duplication, timing, restriction, accessibility, intermediary structures, and long-term funding patterns remain difficult to answer consistently. Capital often appears as isolated transactions rather than as part of a larger landscape of investment and resource movement.

The fourth category is **community-level signals**.
Community experience is increasingly recognized as essential to effective decision-making. Yet community knowledge often enters institutional systems in fragmented ways: surveys, listening sessions, focus groups, stakeholder interviews, community advisory boards, social media, public testimony, and organizational reporting. Even when these signals exist, questions remain about authority, consent, provenance, interpretation, and governance. Community information is not simply another dataset. It carries obligations and relationships that many existing systems were never designed to manage.

The fifth category is **contextual factors**.
Issues do not exist in isolation. Geography, policy environments, infrastructure, transportation, housing, history, culture, trust, language access, institutional relationships, and local conditions all shape how problems are experienced and how solutions function. Context often appears as background information even though it may be one of the strongest determinants of success or failure.

The sixth category is **organization role and fit**.
This is the category most familiar to philanthropy, yet it appears differently when viewed through an issue-level lens. The question is no longer simply whether an organization is strong, capable, or aligned. The question becomes whether that organization is positioned to play a needed role within a broader response landscape. A strong organization may not be the most important responder. A smaller organization may hold unique trust or access. A public agency may control critical infrastructure. An informal network may provide support that no formal institution can replicate. Role and fit become meaningful only when the broader landscape is visible.

Taken individually, none of these six categories is difficult to imagine.

Taken together, they begin to describe something philanthropy does not yet consistently possess.

- Not another database.
- Not another grantmaking platform.
- Not another organizational directory.
- Not another outcome framework.

What emerges instead is the outline of a decision-preparation environment designed around the issue itself.

That distinction is important because it changes the central question.

Instead of asking:
Which organization should receive support?

The decision-maker can first ask:
What does the issue landscape reveal before organizations are compared?

That is the missing layer.

And once that layer becomes visible, a second question emerges.
If such a layer is needed, what kind of infrastructure could support it responsibly?

3. Why the Missing Layer Matters

The absence of a governed issue-level preparation layer matters most when philanthropy moves beyond routine grant administration and into questions of strategy, systems change, place-based investment, community priorities, or long-term issue outcomes. In those situations, funders are not simply choosing between organizations. They are making judgments about the issue itself: which needs are most significant, which responses are already present, which gaps remain, which communities should influence priorities, and which forms of evidence should carry weight. Existing systems can contribute to those judgments, but they rarely assemble the full view. A grants database may show where funding has gone without showing whether it aligns with current need. A nonprofit profile may show organizational capacity without showing whether the organization is positioned to fill the most important role in the ecosystem. A community engagement process may reveal lived experience without connecting those signals to capital flows, institutional response, or policy conditions. The problem is not that these sources are wrong. The problem is that they were not designed to become, together, a governed issue-level decision environment.

This distinction is easy to miss because many philanthropic decisions can proceed without a full issue-level view. A funder renewing support for a long-standing grantee may reasonably rely on organization-level knowledge, relationship history, financial review, and prior reporting. A donor making a values-aligned gift may not need a comprehensive landscape analysis before acting. A grant program with a narrow and well-understood purpose may be able to make responsible decisions using eligibility review, application scoring, and due diligence. The missing layer becomes more consequential when the decision is not merely whether a specific organization deserves support, but whether resources are being directed toward the right roles, in the right places, at the right time, with the right understanding of the issue. At that point, the decision has moved from organizational selection into issue-level judgment.

The research supporting this paper suggests that fragmentation is real, but it is not the deepest problem by itself. Philanthropy does navigate separate systems for nonprofit profiles, grant histories, evidence, community feedback, and capital-flow data. Yet the stronger evidence points to deeper failures involving power asymmetry, capital-flow misalignment, uneven evidence regimes, and accountability gaps. Fragmentation often appears as the visible symptom of these deeper conditions because information is collected for compliance, fundraising, reputation, reporting, or internal administration rather than for shared issue-level understanding. When information is produced for different institutional purposes, it can become difficult to reuse responsibly for issue-governed decision preparation.

The distinction between fragmentation and governance changes the kind of solution that seems appropriate. If the problem were simply that information sits in too many places, the solution might be a larger database or a better search tool. If the problem were simply that grantmaking workflows are inefficient, the solution might be better software. If the problem were simply that funders lack evidence, the solution might be more research. Each of those responses may help in a limited way, but none is sufficient if the deeper problem is that philanthropy lacks a shared, governed way to assemble, interpret, challenge, and use issue-level knowledge before decisions are made.

This is especially important when community-originated information enters the decision environment. Many philanthropic institutions now recognize that communities closest to an issue often hold knowledge that institutional systems cannot easily see. That recognition has led to more listening sessions, community surveys, advisory groups, participatory processes, stakeholder interviews, and lived-experience initiatives. Those practices can improve decision-making, but they also raise governance questions that many existing systems are not prepared to answer. Community knowledge is not simply another input to be collected, tagged, summarized, and stored. It may carry obligations related to authority, consent, provenance, interpretation, benefit, protection, and reuse.

The community data governance research is clear on this point. Across Indigenous data sovereignty, CARE, OCAP, participatory data stewardship, and community-based participatory research, community-originated information is not treated as a neutral institutional asset that becomes freely usable once collected. Communities may be rightsholders, governance participants, or legitimate authorities over collection, access, use, reuse, interpretation, and benefit-sharing, especially where information concerns collective identity, territory, culture, wellbeing, or protected knowledge. For philanthropy, this means community signals become decision-grade only when governance questions have been resolved: who has authority, what consent exists, what restrictions travel with the information, what interpretation has been community-reviewed, and what uses are out of bounds.

This is where the missing layer becomes more than a technical problem. A technical system can connect datasets, normalize fields, create dashboards, generate summaries, and make information easier to retrieve. Those capabilities are useful, but they do not determine whether the resulting issue view is legitimate. Legitimacy depends on whether the system preserves provenance, protects sensitive knowledge, distinguishes community interpretation from institutional analysis, allows correction and challenge, and makes clear who had authority over each part of the decision preparation process. Without those functions, a more integrated system could still reproduce the same extractive or top-down patterns it was meant to correct.

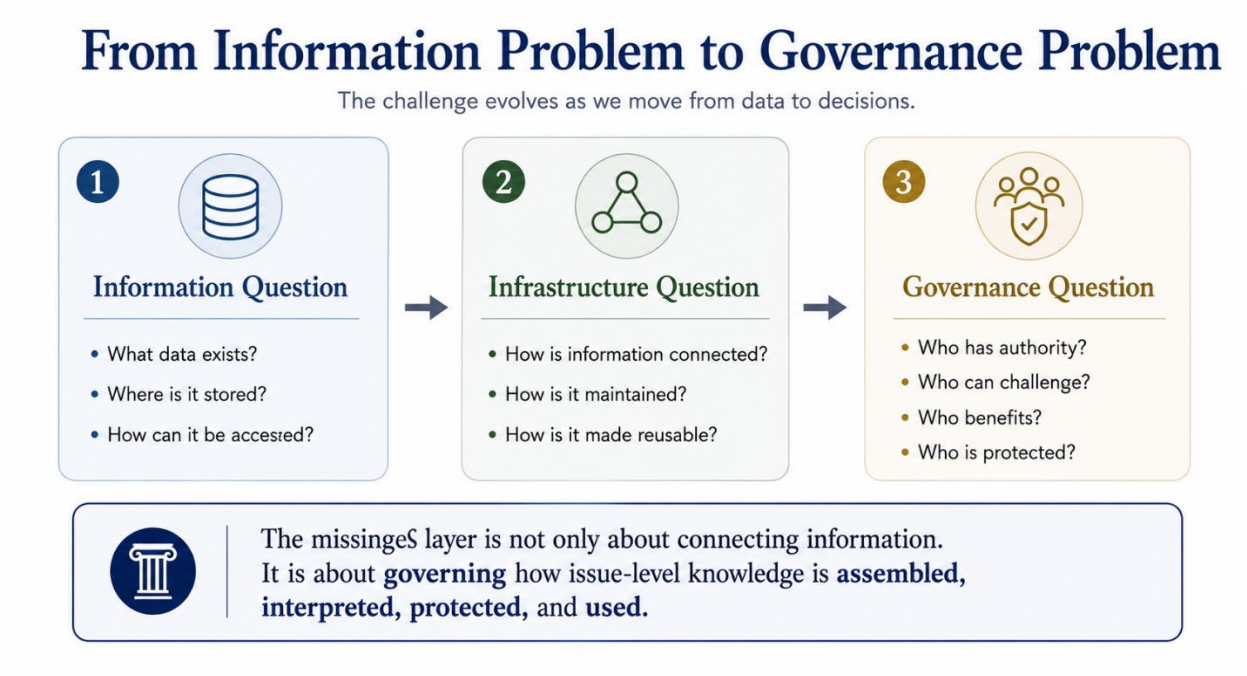

The same concern applies to capital flows. Philanthropy can often identify individual grants, but issue-level preparation requires a different kind of capital visibility. It requires understanding where money is concentrated, where it is absent, whether it is short-term or durable, whether it is restricted or flexible, whether it moves through intermediaries, and whether timing aligns with need. Research on donor-advised funds illustrates why this matters. The strongest evidence does not support a simple claim that DAFs are merely warehouses, but it does show that delay can be real, sometimes substantial, and poorly measured by sponsor-level averages alone. For issue-level preparation, that means it is not enough to know that money exists or that grants eventually occur. Decision-makers need to understand how capital moves in relation to the issue landscape.

The importance of the missing layer also becomes clear when considering large gifts and major philanthropic investments. Evidence on large unrestricted gifts suggests that such gifts often strengthen recipient organizations in the short and medium term, improving financial sustainability, leadership confidence, fundraising capacity, and operational resilience. Yet the evidence becomes thinner as the unit of analysis widens from the recipient organization to the broader ecosystem, field, issue, or community. A gift can be highly beneficial at the organizational level while its ecosystem-level effects remain underexamined. That is precisely the kind of distinction an issue-level preparation layer would help make visible.

In this sense, the missing layer matters because it changes what philanthropy can be accountable to. Without issue-level preparation, accountability often flows through organizational reports, grant compliance, board approvals, and funder-defined outcomes. Those forms of accountability are important, but they do not necessarily answer whether

the issue landscape has changed, whether missing responders became visible, whether capital moved toward the most significant gaps, or whether communities closest to the issue experienced the response as legitimate and useful. A governed issue-level preparation layer would not solve those questions by itself, but it would make them harder to ignore.

The central point is not that every philanthropic decision requires a full infrastructure-backed landscape review. That would be impractical and unnecessary. The point is that philanthropy currently lacks a durable way to support such preparation when it does matter. When decisions become strategic, place-based, collaborative, public-interest oriented, or systems-focused, the sector needs more than scattered information and internal funder interpretation. It needs a governed environment for seeing the issue before selecting organizations, and for documenting how that view was assembled, challenged, protected, and used.

That realization leads to the next question. If philanthropy needs this missing layer, what makes it public-interest infrastructure rather than simply another platform?

4. What Makes Infrastructure Public-Interest

The phrase *public-interest infrastructure* can sound reassuring before it is defined. It suggests something open, useful, civic, and broadly beneficial. But the research reviewed for this paper makes clear that public-interest infrastructure is not simply a label that can be attached to any tool with a social purpose. Nor is it guaranteed by nonprofit status, public funding, open-source code, government involvement, or the availability of a public-facing website. Those features may support a public-interest model, but none of them is sufficient on its own. The more important question is how the infrastructure is governed, who it serves, who can shape it, what rules constrain its use, and whether those rules protect public value over time. Across civic technology, data stewardship, public health, journalism, digital public infrastructure, and nonprofit infrastructure, the common thread is governance and stewardship rather than technology type.

This distinction matters for philanthropy because the missing layer described in this paper would not be neutral. Any infrastructure that helps assemble issue-level knowledge before resources move would influence what becomes visible, whose claims are treated as credible, how capital flows are interpreted, how community signals are handled, and how organizations are understood within a broader landscape. A system with that kind of influence cannot be evaluated only by whether it is useful or efficient. It must also be evaluated by whether its governance arrangements prevent exclusive control over essential functions, data, standards, interfaces, and interpretive authority. The strongest cross-sector conclusion from the public-interest infrastructure research is that public-interest infrastructure exists to create broad public value through accountable governance, inclusive access rules, interoperability, transparency, and long-term stewardship.

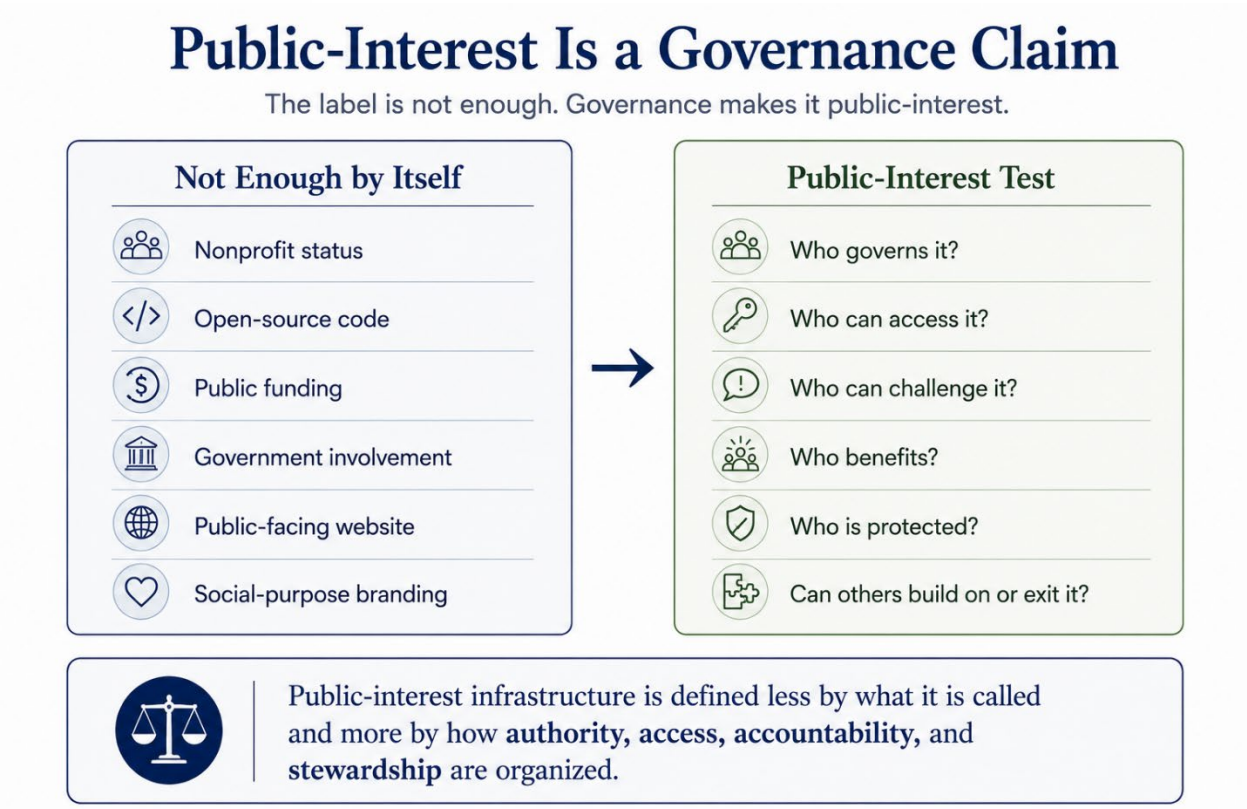

A useful way to see the difference is to compare public-interest infrastructure with platform infrastructure. Platforms can create enormous value, and many proprietary systems in philanthropy already serve important functions. The issue is not that platforms are inherently harmful. The issue is that platforms are usually governed by the owner or operator, which means participation rules, access terms, interfaces, pricing, integrations, and data-use conditions can often be changed according to institutional or commercial priorities. In ordinary software markets, that may be expected. But when a system begins to function as a shared layer for understanding public problems or coordinating resources, unilateral control becomes more consequential. Public-interest infrastructure is designed to reduce that risk by making the rules, standards, access conditions, and stewardship responsibilities more transparent, durable, interoperable, and accountable.

The public-interest infrastructure research points to several examples from other sectors that make this distinction concrete. OpenStreetMap pairs free and open geospatial data with a nonprofit foundation and membership governance. 360Giving provides an open grants-data standard governed through stewardship processes rather than a closed grantmaking environment. UK Biobank uses governed, restricted access tied to public-interest research, ethics review, and committee oversight. DHIS2 and OpenMRS demonstrate how open-source, standards-based health platforms can support local adaptation and integration rather than closed vendor dependence. These examples are not identical, but they share a common logic: the infrastructure becomes valuable because

many actors can rely on it, reuse it, and participate in it without handing exclusive control to a single gatekeeper.

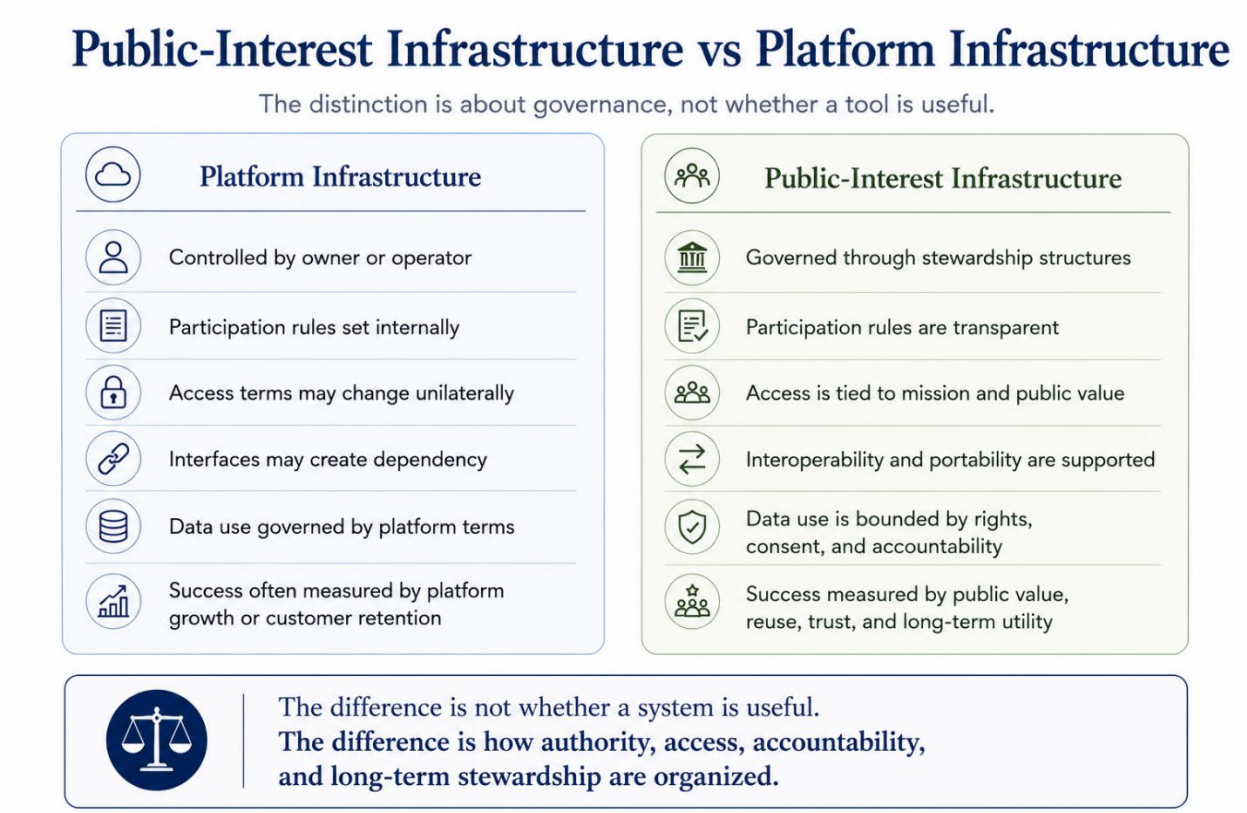

For philanthropy, this distinction is especially important because the missing issue-level layer would sit near the point where knowledge becomes action. It would not merely store information. It would help shape how decision-makers understand issues before allocating resources. If that layer were proprietary, opaque, funder-controlled, or dependent on closed taxonomies and inaccessible data, it could reproduce the same problems Paper 1 identified in organization-first decision-making. It might make issue-level preparation faster or more polished while still leaving fundamental questions about authority, participation, challenge, and protection unresolved. The lesson from adjacent fields is that infrastructure serving public decisions must be designed around public value from the beginning, not retrofitted with public-interest language after control has already been concentrated.

This does not mean that public-interest infrastructure must be free, government-run, or entirely open. The research is careful on this point. Paid services can be compatible with public-interest infrastructure when revenue supports stewardship, maintenance, documentation, security, user support, and long-term sustainability. Restricted access can also be compatible with public-interest infrastructure when information is sensitive, community-originated, legally constrained, or capable of causing harm if released without governance. The question is not whether everything is open or free. The question is whether the rules of access are transparent, mission-aligned, accountable, and designed to prevent

exclusion, enclosure, or unilateral control over an essential layer.

This is where the six recurring characteristics of public-interest infrastructure become useful: governance, access, transparency, accountability, interoperability, and stewardship. These should not be treated as a checklist that can be satisfied by a policy page. They are interdependent conditions. Governance without transparency becomes difficult to scrutinize. Access without accountability can create harm. Interoperability without stewardship can produce brittle systems that no one maintains. Transparency without challenge rights can become performative disclosure. Stewardship without participation can become paternalistic control. The strength of public-interest infrastructure lies in how these characteristics work together to keep essential functions usable, governable, and resistant to capture over time.

The risks that make these characteristics necessary are not hypothetical. Vendor capture can occur when essential functions become dependent on one provider whose terms, pricing, interfaces, or strategic direction can shift over time. Funder capture can occur when philanthropic support builds infrastructure that quietly reflects the priorities of its sponsors rather than the needs of the field. Data enclosure can occur when information of broad public value technically exists but is practically inaccessible, locked behind restrictive terms, incomplete APIs, or proprietary formats. Platform dependency can occur when users become accustomed to convenience inside a closed system while losing the ability to move, audit, reuse, or contest the underlying infrastructure. These risks appear across public administration, journalism, data governance, and civic technology, and they are directly relevant to philanthropic decision preparation.

The implication for Paper 2 is straightforward but demanding. The missing issue-level preparation layer should not be described merely as a new platform. A platform may be the visible interface, but the public-interest claim depends on the deeper structure beneath it. The infrastructure would need transparent governance, stable identifiers, interoperable standards, documented provenance, challenge pathways, access tiers, protections for sensitive and community-originated information, maintenance commitments, and a credible plan for durability beyond a single funder, vendor, or grant cycle. Without those conditions, the sector might gain another useful tool while still failing to build the shared layer required for issue-governed decision preparation.

This is why the next section turns from public-interest infrastructure in general to community-originated information specifically. If issue-governed philanthropy requires community signals, local knowledge, lived experience, and community interpretation, then the infrastructure cannot treat that knowledge as ordinary data. It must be governed differently from the moment it enters the system.

5. Community-Originated Information and Governance

The moment philanthropic infrastructure begins to incorporate community knowledge, the nature of the system changes. A grant record, an IRS filing, a public dataset, a program budget, and a census indicator each carry their own limitations, but they are generally institutional or administrative records. Community-originated information is different. It may include lived experience, local interpretation, cultural knowledge, trust signals, community priorities, informal warnings, stories of harm, descriptions of barriers, or knowledge about relationships and places that do not appear in official systems. Those forms of knowledge can be essential to understanding an issue, but they are not neutral inputs that become freely usable once collected.

This distinction matters because issue-governed philanthropy cannot rely only on institutional data. If the goal is to understand what an issue requires before resources move, then community experience must have a place in the decision-preparation environment. Communities often know where formal systems fail, which services are

trusted, which barriers are understated, which organizations have legitimacy, and which solutions do not fit local reality. Yet the value of community knowledge does not give philanthropy the right to extract it, summarize it, store it, or reuse it without governance. The more important community-originated information becomes to decision-making, the more carefully its authority, consent, provenance, interpretation, and protection must be handled.

The community data governance research reviewed for this paper is clear on this point. Across Indigenous data sovereignty, CARE, OCAP, AIATSIS research ethics, participatory data stewardship, and community-based participatory research, a common principle emerges: community-originated information is not simply an institutional asset. Communities may be rightsholders, governance participants, or legitimate authorities over collection, access, use, reuse, interpretation, and benefit-sharing, especially when information concerns collective identity, territory, culture, wellbeing, protected knowledge, or community experience. External institutions may serve as stewards, but stewardship is not ownership.

That principle creates a very different standard for philanthropic decision preparation. It is not enough to ask whether community signals have been collected. It is not enough to ask whether community members were consulted, surveyed, interviewed, or invited to participate in a listening session. Those activities may be useful, but they do not resolve the deeper governance questions. Who had authority to authorize collection? What consent was given? What uses were permitted? What restrictions travel with the information? Who interpreted the meaning of what was shared? What should remain protected, attributed, restricted, or excluded from decision systems altogether?

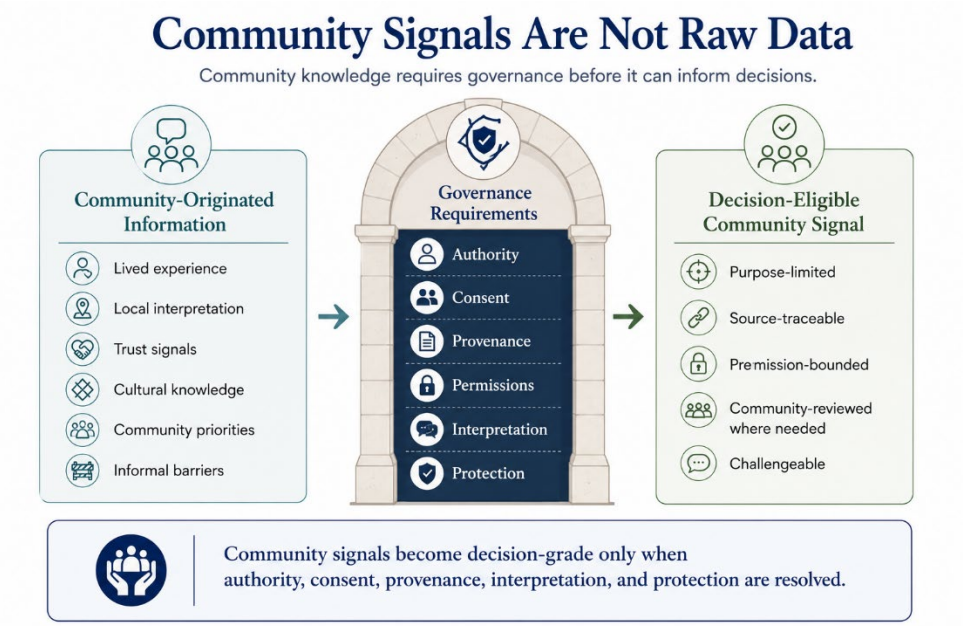

The difference between collection and governance is especially important because philanthropy has often treated community input as an ingredient that can be added to an otherwise funder-controlled process. A funder may commission interviews, hold focus groups, convene community advisors, or collect stories from affected people and then fold those insights into an internal memo, strategic plan, or grantmaking rubric. That may improve the quality of the funder's understanding, but it does not necessarily change who controls the meaning of the information. Community knowledge can still be filtered through institutional categories, interpreted against funder-defined priorities, and used to justify decisions that communities had little authority to shape.

This is why the governance frameworks are so useful for philanthropic infrastructure. CARE shifts attention from data availability toward collective benefit, authority to control, responsibility, and ethics. OCAP frames ownership, control, access, and possession as concrete governance rights for First Nations data. Participatory data stewardship emphasizes participation across the data lifecycle rather than one-time consultation. Community-based participatory research treats community stakeholders as partners in research design, collection, analysis, interpretation, dissemination, and sustainability. These frameworks come from different histories and contexts, and they should not be collapsed into one generic model. Still, together they offer a clear warning: information about communities should not be governed only by the institutions that collect it.

One of the most important implications is that community interpretation and institutional interpretation should be kept distinct. A foundation, consultant, platform, or intermediary may summarize what it heard from community members, but that summary is not the same as community meaning. Institutional interpretation can be useful, but it should be labeled as such. Community-reviewed interpretation carries a different authority because it has been examined, corrected, contextualized, or affirmed by people with standing in the community. Local Contexts' Community Voice approach and community-based participatory research both reinforce this distinction by recognizing that existing institutional descriptions may be incomplete, partial, or disconnected from community meaning.

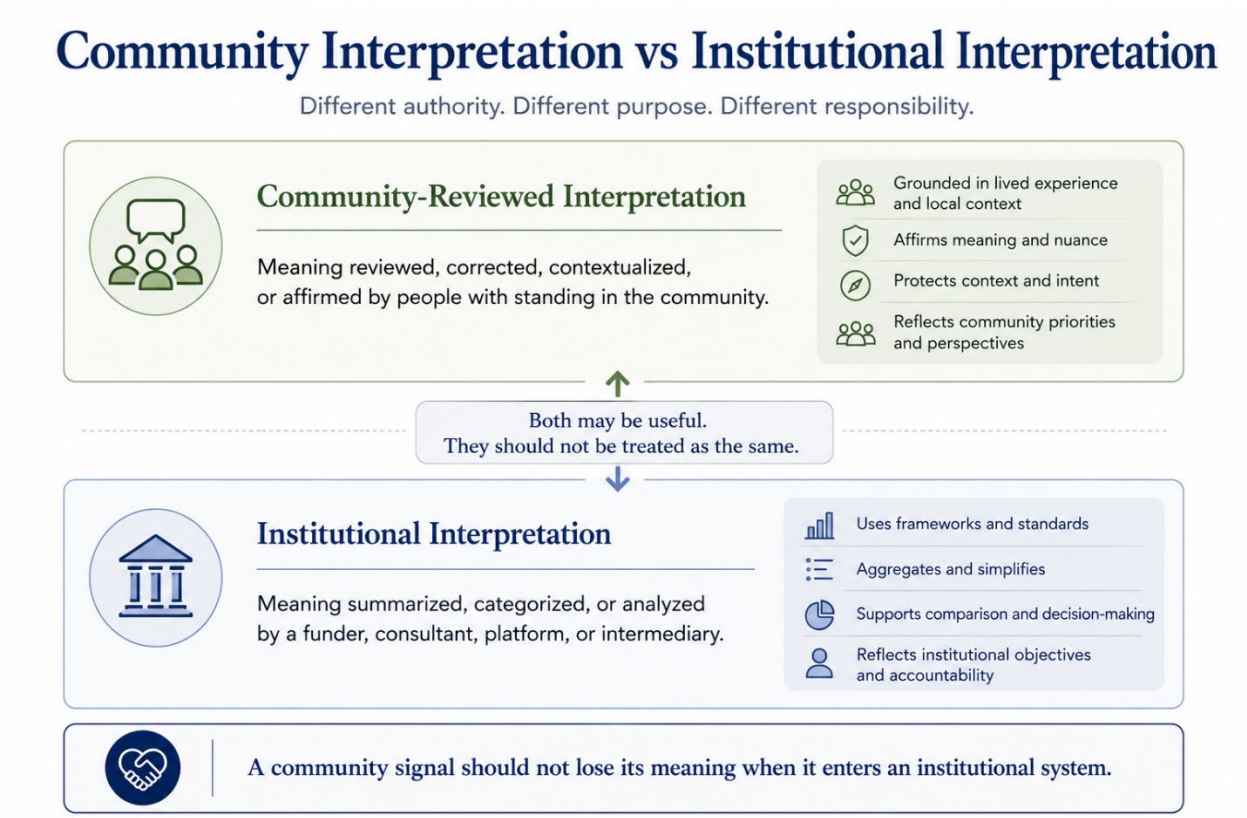

The same care is needed with provenance. In conventional data systems, provenance often means documenting where a fact came from, when it was retrieved, and how it was transformed. That is necessary, but community-originated information often requires more than technical provenance. It may require relational and cultural provenance as well: who shared the information, under what relationship, with what authority, for what purpose, under what permissions, with what restrictions, and with what expectation of review or benefit. The community governance research recommends documenting both the technical chain and the relational context because the meaning and permitted use of community information cannot be separated from the conditions under which it was shared.

This has direct consequences for infrastructure design. A system that merely stores community feedback in a text field is not sufficient. Nor is a dashboard that aggregates community sentiment without preserving context, authority, or restrictions. If community-originated information is going to inform philanthropic decisions, the infrastructure must be able to carry permissions, consent status, source community, interpretation status, sensitivity flags, attribution requirements, review rights, permitted uses, prohibited uses, and retention limits. The governance layer cannot be a policy document sitting outside the system. It has to travel with the information itself.

Some information should not enter philanthropic decision systems at all. This is a hard but necessary point. The community governance research identifies categories of knowledge that may need to be protected, restricted, or excluded unless a community affirmatively authorizes use. These include sacred or secret knowledge, ceremonial materials, gender-restricted or seasonally restricted knowledge, sacred-site locations, culturally sensitive information, commercially valuable traditional knowledge, and information whose circulation could deepen trauma, stigma, cultural loss, or community risk. Anonymization may reduce some individual privacy risks, but it does not resolve collective harms when information concerns a people, place, culture, or community relationship.

That point is particularly important for issue-governed philanthropy because the desire for better issue visibility can unintentionally create new forms of extraction. A funder may want to understand community barriers more fully. A collaborative may want to identify trusted local responders. A platform may want to synthesize lived-experience signals across multiple sources. Each of those goals may be legitimate, but each can also create risk if community knowledge is removed from its context, made searchable beyond its intended audience, translated into indicators without review, or reused for purposes that were never authorized. The infrastructure should therefore begin from a presumption of stewardship rather than institutional ownership.

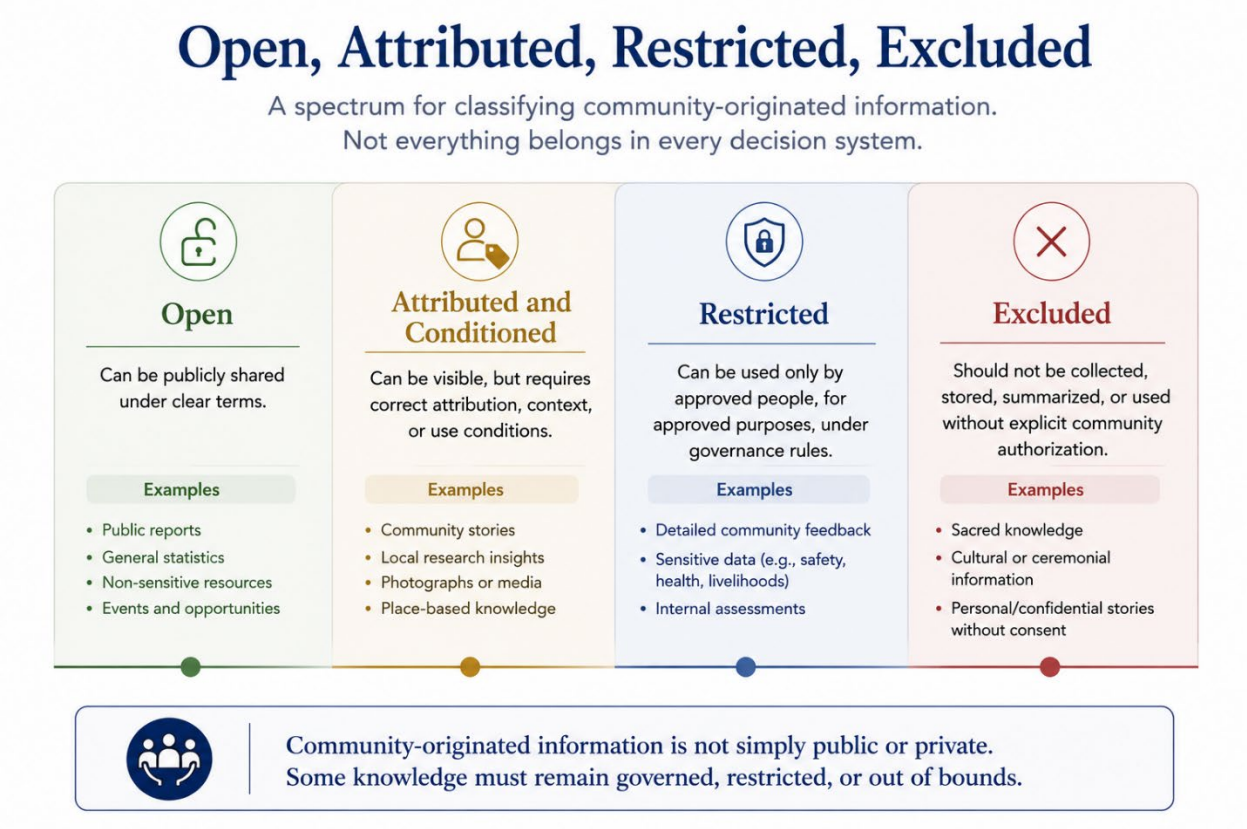

The safeguards required here are not peripheral to Paper 2. They are central to the meaning of public-interest infrastructure. If the missing issue-level layer is meant to support more legitimate decision preparation, it must not turn community knowledge into another extractive resource. At minimum, the infrastructure should include an authority map, a consent and permissions register, provenance records that travel with the information, sensitivity and exclusion rules, separation between community interpretation and institutional interpretation, benefit and feedback loops, and funding for governance capacity. These are not merely ethical ideals. They are operating requirements if community signals are going to be used responsibly in philanthropic decisions.

The authority map is especially important because the word community can hide complexity. Authority may sit at different levels depending on the information involved. Some knowledge may be held by a nation, tribe, or community body. Some may be held by a subgroup, family, clan, cultural authority, organization, or named individual. Some may be appropriate for open community discussion, while other knowledge may be restricted by protocol. A philanthropic institution cannot assume that one participant, one nonprofit, one consultant, or one advisory group has authority to approve all uses of community-originated information.

Consent should also be understood as ongoing and purpose-specific rather than one-time and universal. A person or community may agree to share information for one purpose and not another. They may permit a story to inform a local strategy discussion but not a public dashboard. They may allow a theme to be summarized but not allow raw notes to be stored. They may approve use for one decision cycle and require review before reuse. Infrastructure that cannot represent these distinctions will push community knowledge into institutional categories that are too blunt for responsible use.

The same is true for benefit and feedback. Communities should not be asked to provide knowledge into systems that never return anything useful to them. A public-interest issue-preparation layer should be designed so that communities and responders can see how information was used, correct errors, challenge interpretations, receive usable outputs, and understand what changed as a result. Otherwise, even well-intentioned listening becomes a one-way transfer of value from communities to institutions.

The practical implication is clear. Philanthropic institutions should describe themselves as delegated stewards, not owners, of community-originated information. They should define community signals as governance-bound inputs rather than neutral data. They should treat information as eligible for decision use only when it is authority-validated, purpose-limited, provenance-documented, permissions-bounded, and paired with community-reviewed interpretation where needed. Information that fails those tests should be treated as advisory only, heavily restricted, or excluded entirely.

This does not make issue-level preparation impossible. It makes it more honest. It recognizes that better visibility is not the same as legitimate use. It also prevents the missing infrastructure layer from solving one problem while creating another. If philanthropy wants to understand issues before resources move, it must build infrastructure that can protect the knowledge needed for that understanding.

The next section turns from these governance principles to the practical infrastructure question: what is the smallest credible system that could support governed issue-level decision preparation without becoming a giant centralized data warehouse?

6. Minimum Viable Infrastructure for Issue-Governed Decision Preparation

Once the need for a governed issue-level preparation layer is established, the conversation naturally shifts from principles to implementation. At that point, an important risk emerges. Readers may begin imagining a massive new system: a giant database containing every grant, every organization, every outcome, every community signal, every public dataset, and every piece of philanthropic knowledge in one place. The history of information systems suggests caution. Large centralized visions are often attractive on paper but difficult to govern, expensive to maintain, slow to adapt, and vulnerable to both technical and institutional failure.

The research supporting this paper points toward a different conclusion. The smallest credible infrastructure for issue-governed decision preparation is not a comprehensive warehouse and it is not a universal source of truth. Instead, it is best understood as a relatively thin, federated layer that helps connect existing information while preserving provenance, governance, accountability, and local stewardship. Much of the underlying infrastructure already exists. What remains underdeveloped is the connective tissue that allows issue-level preparation to occur across systems rather than inside isolated silos.

This finding is encouraging because it changes the scale of the challenge. Building an entirely new philanthropic information universe would require enormous resources and would likely reproduce many of the problems discussed earlier in this paper. Building a governed preparation layer is more realistic. It allows organizations, funders, communities, intermediaries, researchers, public agencies, and existing platforms to continue operating within their own environments while contributing to a shared decision-preparation ecosystem.

At its simplest, the infrastructure required for issue-governed preparation can be understood as six connected layers. None of these layers is particularly valuable in isolation. Their importance emerges from the way they work together to create a traceable, challengeable, and governed view of an issue before resources are allocated.

The first layer is a registry and crosswalk layer. Every decision environment depends on the ability to identify what is being discussed. Organizations, places, grants, programs, issues,

populations, and sources all need stable identifiers if information is going to be connected across systems. The research strongly suggests that philanthropy does not need to invent entirely new identifier systems. Existing identifiers already exist for many purposes. The challenge is not creating new registries but maintaining reliable crosswalks among them. The practical task is helping a decision-maker understand that different systems may be describing the same organization, place, issue, or grant in different ways.

The second layer is a grants and transactions layer. Philanthropy is ultimately concerned with the movement of resources, yet many discussions of infrastructure focus almost exclusively on organizations or outcomes. An issue-level preparation environment requires visibility into how capital moves across an issue, where it concentrates, where it is absent, how it changes over time, and how it reaches responders. The research repeatedly points toward grants and transaction records as a foundational element because they provide one of the few ways to see how intentions become actual resource flows.

The third layer is a source catalog. One of the recurring challenges identified throughout this paper is that information is scattered across reports, datasets, evaluations, administrative records, community assessments, public indicators, websites, APIs, and many other forms of evidence. A decision-preparation environment does not need to absorb all of these sources into a single database. It does need a reliable way to describe them, locate them, understand their limitations, and connect them to the claims that emerge from them. A source catalog makes evidence discoverable and reviewable without requiring centralized ownership.

The fourth layer is provenance. Provenance may be one of the least visible but most important parts of the entire infrastructure. When a decision-maker encounters a claim, a summary, a coded observation, a community signal, a capital-flow analysis, or an issue-level interpretation, there must be a way to understand where it came from. Provenance records who created information, who transformed it, what sources contributed to it, what version was used, and what changes occurred along the way. Without provenance, issue preparation quickly becomes vulnerable to misplaced trust, hidden assumptions, and unchallengeable conclusions. The infrastructure research repeatedly identifies provenance as one of the most critical safeguards in any future issue-level system.

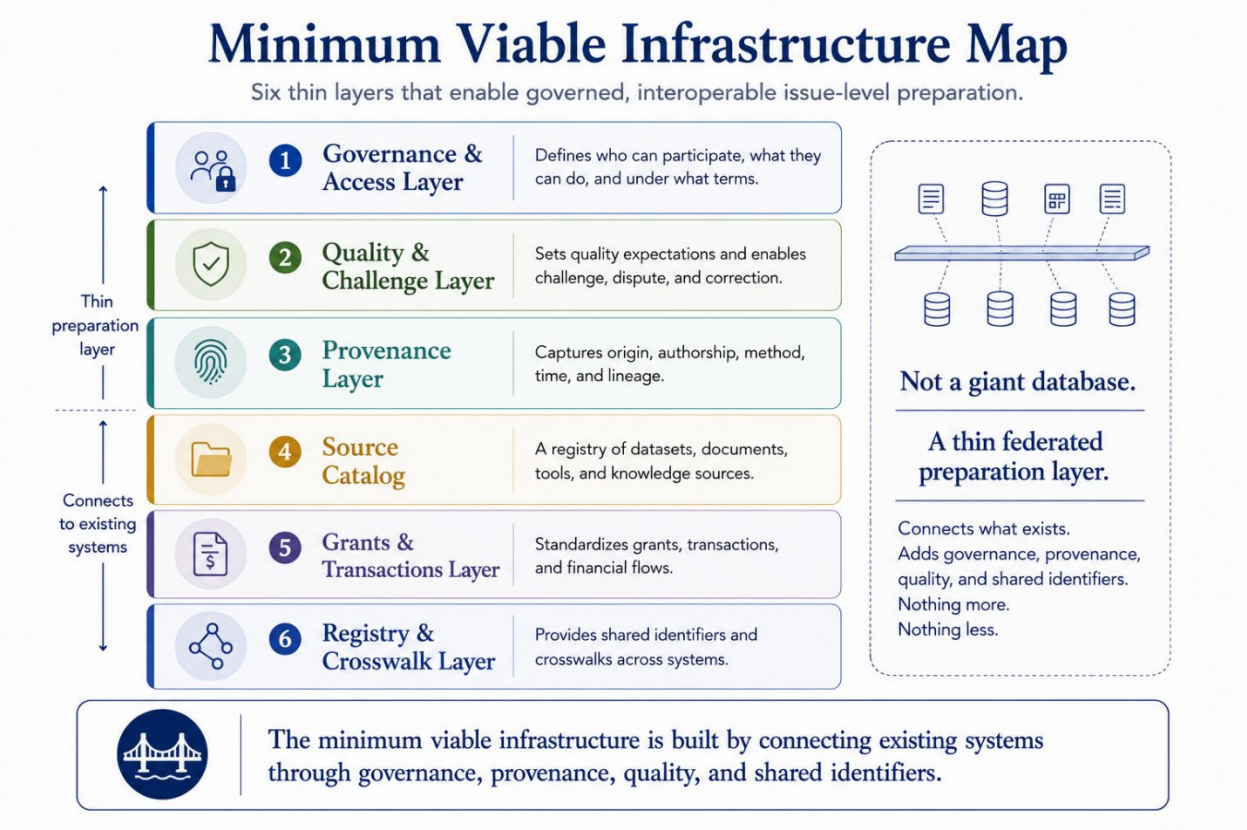

The fifth layer is a quality and challenge layer. Most philanthropic systems contain some form of quality review, but issue-level preparation requires something broader than internal verification. Information should be challengeable. Claims should be reviewable. Contradictions should be visible. Corrections should have pathways. One of the most consistent findings across infrastructure governance research is that quality improves when systems support challenge rather than merely publication. OpenStreetMap, ROR, Wikidata, and other mature infrastructure initiatives all depend on some form of challenge-and-resolution process rather than assuming that information enters the system perfectly formed.

This observation is especially important because issue-level preparation often involves interpretation rather than simple fact collection. Different actors may reasonably disagree about what a pattern means, whether a gap exists, how a community signal should be understood, or which response appears most promising. A challenge layer does not eliminate disagreement. It creates a place where disagreement can be documented, examined, and governed.

The sixth and final layer is governance and access. Earlier sections argued that governance is not an optional feature that can be added later. The infrastructure research arrives at the same conclusion from a different direction. Every significant decision about identifiers,

taxonomies, permissions, access, challenge rights, stewardship responsibilities, transparency obligations, and change control eventually becomes a governance question. The technical architecture can support those decisions, but it cannot make them disappear. Governance determines who can participate, who can challenge, who can authorize changes, who can access different classes of information, and what protections exist for sensitive or community-originated knowledge.

Taken together, these six layers create something more useful than a collection of disconnected systems and more realistic than a centralized philanthropic warehouse. They create a preparation environment. The distinction matters because the purpose of the infrastructure is not to become the destination for all philanthropic information. Its purpose is to help decision-makers assemble a governed understanding of an issue before resources move.

One of the most significant findings in the infrastructure research is that the architecture should be federated rather than centralized. This principle appears repeatedly across successful infrastructure initiatives. Information remains with the organizations, institutions, communities, and systems that produce it. Catalogs point to sources rather than absorbing them. Registries connect identifiers rather than replacing them. Governance coordinates participation rather than demanding uniform ownership. This approach is especially well suited to philanthropy because information comes from many actors operating under different obligations, permissions, sensitivities, and governance structures.

The implications become clearer when considering what the infrastructure does not require. It does not require a universal issue ontology. It does not require a universal population registry. It does not require every organization to use the same software. It does not require a single authority to control all information. The research explicitly cautions against claiming that these harder ontology and governance questions have already been solved. Instead, the evidence supports a more modest and more achievable path: shared identifiers, mapped taxonomies, source-linked evidence, provenance tracking, challenge workflows, and governed access.

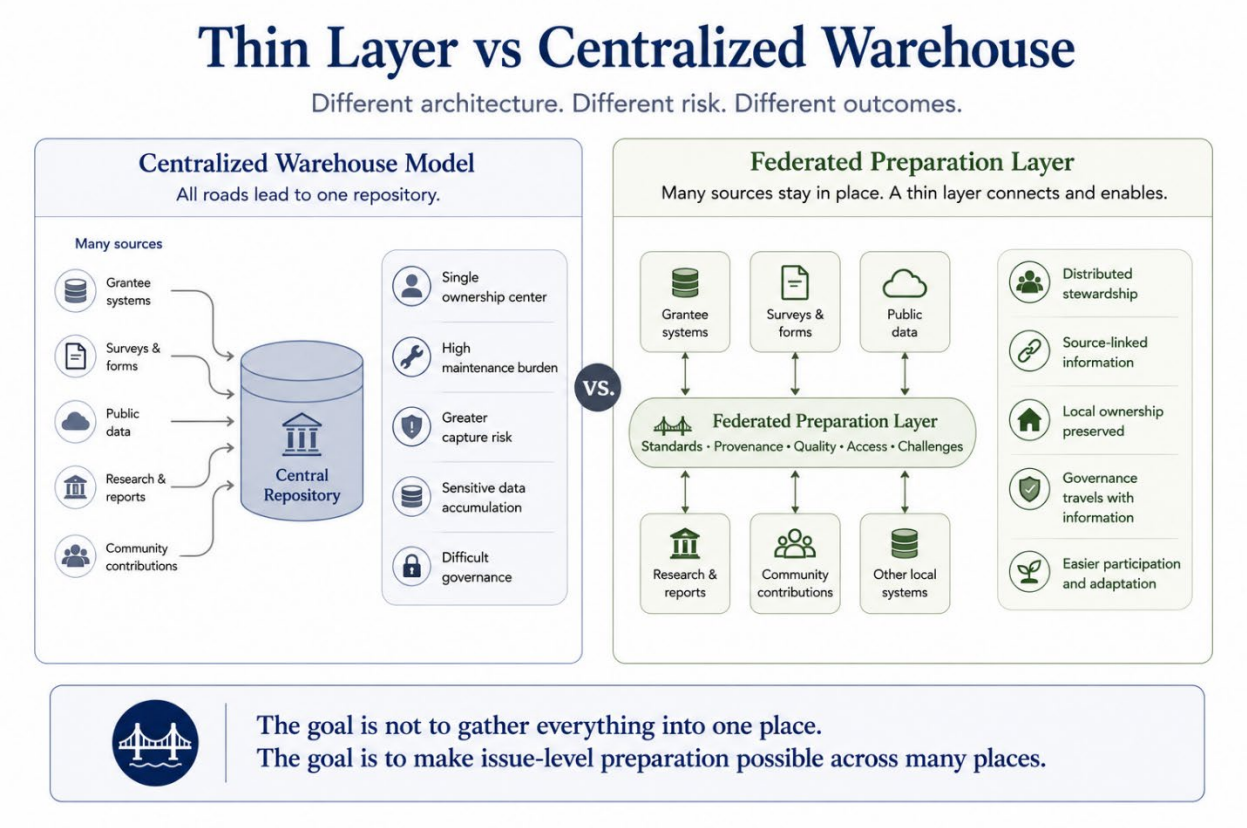

Perhaps the most important lesson from the infrastructure research is that the minimum viable version is smaller than many people imagine. The infrastructure does not need predictive scoring engines, universal taxonomies, automated recommendations, or a comprehensive database of every philanthropic activity. It needs enough structure to support governed issue preparation. That means stable identifiers, visible capital flows, source-linked evidence, documented provenance, challenge mechanisms, and governance capable of managing participation, authority, and protection. Everything else can evolve over time.

This is an important conclusion because it transforms the problem from an impossible ambition into a practical design challenge. The missing layer is not missing because the technology is unattainable. It is missing because governance, stewardship, and issue-level preparation have historically received less attention than grantmaking itself.

The next section explores what happens when these infrastructure principles are combined with the governance principles introduced earlier. It introduces a model that is increasingly common across public-interest infrastructure efforts: an open core paired with a governed edge.

7. Open Core, Governed Edge

One of the recurring tensions in discussions about philanthropic infrastructure is the apparent conflict between openness and protection. Advocates of transparency often argue that more information should be publicly available. Advocates of privacy, community governance, and responsible stewardship point to the risks of unrestricted access, especially when information concerns vulnerable populations, culturally sensitive knowledge, community narratives, or protected forms of expertise. Viewed as a simple choice between open and closed systems, the tension can appear difficult to resolve.

The research reviewed for this paper suggests that the dilemma is often framed too narrowly.

The most successful examples of public-interest infrastructure rarely operate at either extreme. They are neither fully open nor fully closed. Instead, they combine open layers that support transparency, interoperability, discovery, and accountability with governed layers that protect sensitive information, preserve rights, and respect legitimate authority. In other words, they separate what should be broadly visible from what should remain governed. This pattern appears across digital public infrastructure, health information systems, open data initiatives, scientific repositories, grants-data standards, and community-governance frameworks. The underlying principle is surprisingly consistent: openness should apply where openness creates public value, while governance should apply where unrestricted access creates risk.

For issue-governed philanthropy, this distinction is especially important because the infrastructure envisioned in this paper must simultaneously support visibility and protection. Decision-makers need enough openness to understand who is acting, where resources are flowing, what evidence exists, how claims are supported, and how issue landscapes are changing over time. At the same time, they need enough governance to protect community-originated information, sensitive records, restricted knowledge, private data, and forms of information that could cause harm if broadly distributed.

The research points toward what might be called an **open core, governed edge** model. The idea is straightforward. Some parts of the infrastructure should be openly accessible because they become more valuable when many people can discover, examine, reuse, and challenge them. Other parts should remain governed because their value depends on responsible stewardship rather than unrestricted circulation. The challenge is not deciding whether the system is open or closed. The challenge is deciding which information belongs in which layer and what rules should govern movement between them.

The open core contains the shared foundations that make issue-level preparation possible across institutions. Registry information, common identifiers, public metadata, source descriptions, governance policies, taxonomies, provenance structures, quality indicators,

challenge mechanisms, and many organization-level grant records can often be made broadly accessible. These elements help create transparency, reduce duplication, improve interoperability, and support external scrutiny. They allow multiple actors to participate in a shared ecosystem without requiring dependence on a single platform or institution. The infrastructure research repeatedly points toward open metadata, open registries, open provenance structures, and reusable standards as foundational components of durable public-interest infrastructure.

The governed edge serves a different purpose. It exists to protect information that cannot be responsibly managed through openness alone. Community-originated information may carry permissions, restrictions, cultural protocols, review requirements, or collective rights. Individual-level records may involve privacy concerns. Certain forms of local knowledge may be sensitive, contextual, or vulnerable to misuse if removed from their original setting. Some information may be legally restricted. Other information may be ethically restricted even when legal barriers do not exist. The governed edge acknowledges that responsible stewardship sometimes requires controlled access, review processes, role-based permissions, and ongoing accountability.

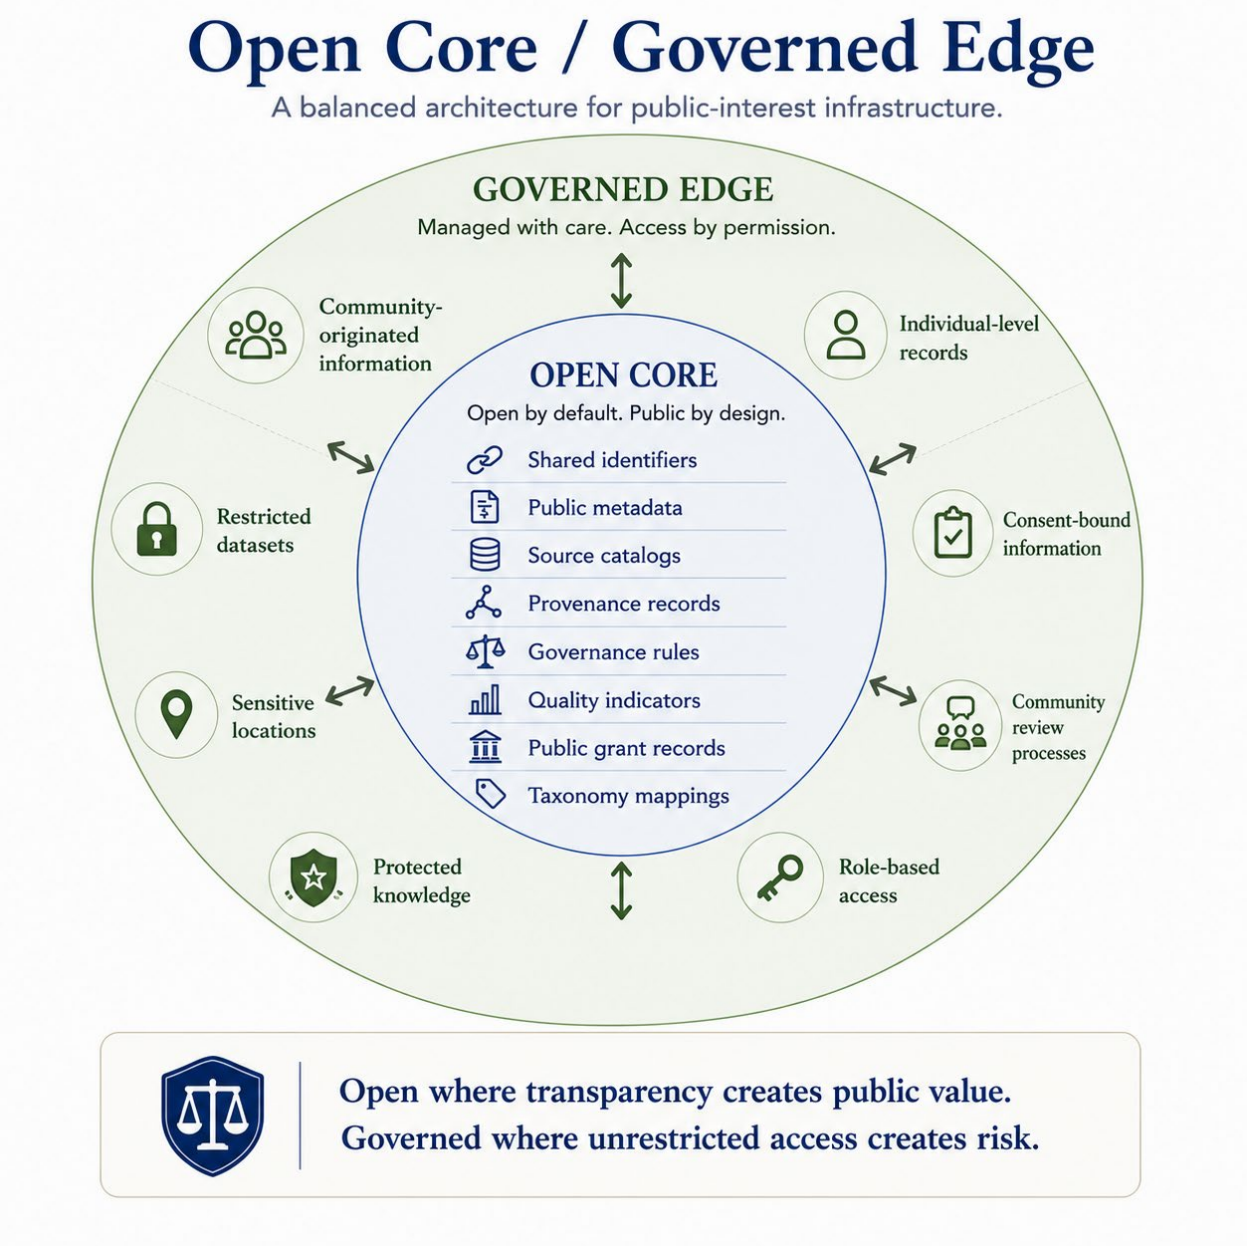

One of the advantages of this model is that it reduces the tendency to treat every information governance decision as an all-or-nothing debate. Too often, infrastructure discussions become trapped between two unsatisfactory positions. One position assumes that openness is inherently good and that information should be made available whenever possible. The other assumes that sensitivity requires information to remain inaccessible except within tightly controlled institutional environments. Neither approach is well suited to the complexity of issue-governed philanthropy. Some information becomes more valuable when widely shared. Other information becomes more dangerous when widely shared. The infrastructure therefore needs mechanisms capable of distinguishing between these cases rather than forcing them into the same category.

This distinction also helps clarify how community governance fits into the broader architecture. Community-originated information does not become part of the open core simply because it is useful. The governance research repeatedly emphasizes that authority, consent, provenance, interpretation, and benefit-sharing remain attached to community knowledge even after it enters institutional systems. The open core may contain descriptions that certain information exists, metadata about how it is governed, or references to the existence of community-reviewed interpretations. The underlying information itself may remain within governed spaces controlled by permissions, review processes, and community authority structures.

The same logic applies to issue preparation more broadly. A funder may need to know that evidence exists regarding a community concern, a trust barrier, a culturally specific service model, or a local response network. That does not necessarily mean the funder needs unrestricted access to every underlying record. Public-interest infrastructure often works by making information discoverable and governable at the same time. Discovery allows decision-makers to know what exists. Governance determines how it may be accessed, interpreted, or used.

The open core and governed edge model also provides a practical answer to one of the most common concerns raised about issue-level infrastructure: the fear that it will eventually become another centralized repository. Centralization tends to create pressure toward uniform rules, concentrated authority, and increasing sensitivity risks as more information accumulates in one place. A federated model reduces those pressures. Information can remain close to its original stewards while still participating in a shared decision-preparation environment. Communities can retain authority over community-originated information. Organizations can maintain control over organizational records. Public agencies can govern public data according to their own requirements. The infrastructure provides connection rather than absorption.

They are different responsibilities that apply to different parts of the same system. Public-interest infrastructure succeeds when it allows these responsibilities to coexist rather than forcing one to dominate the others.

This idea also provides an important transition to the final sections of the paper. Once the architecture becomes visible, the next question is not how to build everything at once. The next question is what claims can responsibly be made about such infrastructure, what risks remain unresolved, and what limitations should be acknowledged before moving from concept to implementation.

8. Limits, Risks, and Responsible Claims

Every infrastructure vision carries a temptation toward overstatement.
Once a pattern becomes visible, it is easy to begin imagining that a sufficiently sophisticated system could solve it. The history of philanthropy, public administration, technology, and social change suggests otherwise. Most enduring problems survive not because information is absent, but because reality is more complicated than any single framework can fully capture.

For that reason, this paper is strongest when its claims remain disciplined.

- The argument is not that philanthropy currently lacks data. It does not.
- The argument is not that philanthropy lacks software. It does not.
- The argument is not that philanthropy lacks intelligence, evidence, expertise, community engagement, evaluation methods, or grantmaking systems.
- Nor does this paper argue that a new infrastructure layer would eliminate difficult judgments, conflicting priorities, political realities, power imbalances, institutional incentives, or uncertainty.

The claim is narrower.

The claim is that philanthropy lacks a durable, governed, issue-level decision-preparation environment capable of bringing together conditions of need, institutional response, capital flows, community signals, contextual factors, and organization role and fit before resources are allocated. Existing systems address portions of this challenge. What remains underdeveloped is the connective layer that allows these dimensions to be understood together rather than separately.

This distinction is important because the paper's credibility depends on resisting claims that the research does not support.

For example, the paper should not suggest that philanthropy currently has no visibility into issue landscapes. That would be inaccurate. Existing systems such as Candid, PolicyMap, Impact Genome, GiveWell, Giving Compass, and others already provide meaningful visibility into organizations, grants, geography, outcomes, interventions, and evidence. In many cases these systems are sophisticated, widely used, and continuously improving. The gap identified throughout this paper is not the absence of visibility. It is the absence of a coherent issue-level preparation environment that brings those forms of visibility together in a governed way before allocation decisions are made.

The paper should also avoid implying that the technical foundations of issue-level preparation are already solved. In several important areas, they are not.

One example is issue classification. The infrastructure research found no universally accepted issue taxonomy capable of spanning the full range of philanthropic activity.

Multiple classification systems already exist, each serving different purposes and operating at different levels of specificity. The practical challenge is not choosing a single winner. The challenge is supporting mappings among systems while acknowledging that classification remains a governed activity rather than a settled technical fact.

The same caution applies to population definitions. Organizations are relatively easy to represent as persistent entities. Populations are different. Age, geography, disability, income, ethnicity, cultural affiliation, and many other characteristics can intersect in ways that resist simple classification. The infrastructure research explicitly warns against assuming that one universal population ontology already exists or can be imposed without consequence. Population description remains a matter of governance, context, and responsible use.

Community-originated information presents an even more complex challenge. Earlier sections argued that community signals are not raw data. The governance research reinforces that point by showing that authority, consent, provenance, interpretation, and benefit-sharing remain attached to community-originated information even after it enters institutional systems. This means there is no responsible path toward simply aggregating all community knowledge into a common repository and treating it as universally available decision support. Any infrastructure that attempts to do so would reproduce the extractive patterns these governance frameworks were designed to prevent.

Throughout the research, one lesson appeared repeatedly across public-interest infrastructure, civic technology, participatory governance, and philanthropic reform. Infrastructure can improve visibility. Infrastructure can improve coordination. Infrastructure can improve traceability. Infrastructure can improve accountability. Infrastructure cannot, by itself, resolve disagreements about values, priorities, tradeoffs, power, or strategy.

This matters because the infrastructure proposed in this paper will inevitably be interpreted by people. Decision-makers will still disagree about which needs are most urgent. Communities will not always speak with a single voice. Evidence may be incomplete, contradictory, or contested. Funders may hold different theories of change. Organizations may disagree about what constitutes success. Infrastructure can create better conditions for those conversations. It cannot eliminate the need for them.
In fact, one of the more important risks may be the opposite. Better infrastructure can create the illusion that difficult judgments have become objective simply because they are supported by more data. This risk appears in many fields. Dashboards, rankings, indicators, scores, maps, and models can create a sense of certainty that exceeds what the underlying evidence supports. The challenge is not merely making information visible. The challenge is making uncertainty visible as well.

For that reason, issue-governed infrastructure should be designed not only to reveal evidence but also to reveal limitations. Decision-makers should be able to see where

information is incomplete, where sources disagree, where community review has not occurred, where classifications are uncertain, where evidence is weak, and where interpretation remains contested. In some cases, one of the most valuable outputs of a decision-preparation system may be a clearer understanding of what is not known.

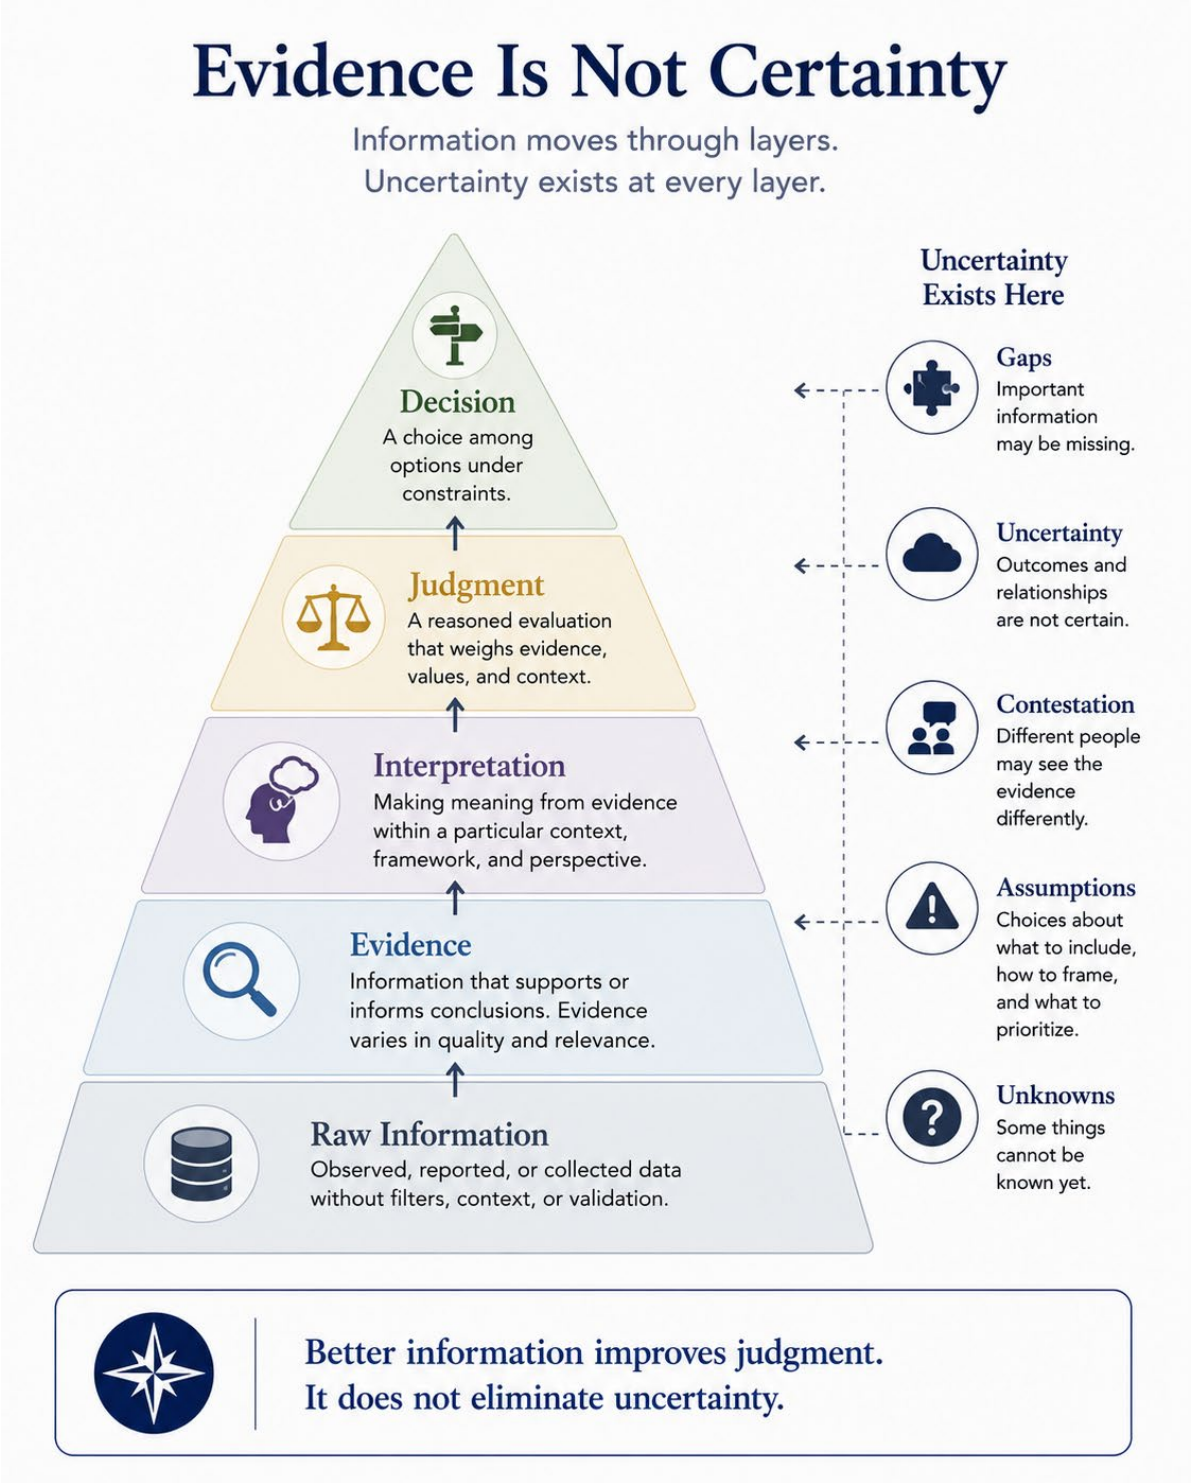

The final limitation concerns scale and adoption. Public-interest infrastructure does not become infrastructure merely because it is designed well. It becomes infrastructure when people rely upon it, contribute to it, govern it, maintain it, challenge it, and build upon it over time. Many technically capable systems fail because governance is weak. Others fail because participation remains limited. Some fail because funding disappears. Others become dependent on a single institution whose priorities eventually change.

The infrastructure research repeatedly emphasizes stewardship as a core requirement rather than an afterthought. Long-term viability depends not only on architecture but on institutions capable of maintaining trust, resolving disputes, adapting standards, managing participation, and surviving leadership transitions. This may be one of the least glamorous aspects of infrastructure, but it is often the most important. Infrastructure succeeds when it remains useful after the original builders are no longer at the center of the story.

Seen in this light, the contribution proposed by this paper is intentionally modest.

- It does not promise perfect decisions.
- It does not promise complete visibility.
- It does not promise consensus.
- It does not promise objectivity.

What it proposes is something more practical: a public-interest infrastructure layer capable of helping philanthropy understand issues more completely, govern knowledge more responsibly, and prepare decisions with greater transparency, accountability, and legitimacy before resources move.

That is a meaningful ambition.

It is also a realistic one.

9. Conclusion: Infrastructure for Seeing Before Selecting

This paper began with a practical question that follows directly from the argument of the first paper in this series. If philanthropy often uses organization-level processes to answer issue-level questions, and if issue-governed decision preparation is necessary, then what would need to exist for that preparation to become possible before resources move?

The answer is not simply more data. Philanthropy already has more information than any single decision-maker can reasonably absorb. Nor is the answer simply more software. The sector already contains serious platforms for nonprofit discovery, grantmaking workflow, outcome analysis, place-based indicators, donor guidance, reporting, and portfolio management. These systems matter, and they have improved many parts of philanthropic practice. They help users find organizations, process applications, track grants, compare outcomes, examine places, and manage portfolios. The problem is not that the sector has failed to build infrastructure. The problem is that the infrastructure it has built is usually

organized around separate pieces of the decision environment rather than around the issue landscape itself.

That distinction is the central contribution of this paper. Existing systems can make organizations visible, grants visible, outcomes visible, places visible, applications visible, and portfolios visible. What remains underdeveloped is a governed public-interest layer that brings these forms of visibility together before allocation decisions are made. Without such a layer, philanthropy may continue to ask issue-level questions while relying on systems organized around organizations, grants, workflows, donor preferences, or historical funding patterns. Each of those views is useful. None is the whole issue.

The missing layer described in this paper is not a giant centralized database or a universal source of truth. In fact, the research points in the opposite direction. A credible minimum infrastructure would be thinner, more federated, and more governable. It would use registries, crosswalks, transaction records, source catalogs, provenance systems, quality and challenge workflows, access rules, and stewardship structures to help decision-makers assemble issue-level knowledge across existing systems. Its purpose would not be to absorb every piece of philanthropic information into one place. Its purpose would be to make issue-level preparation possible across many places.

This is why the concept of public-interest infrastructure matters. The missing layer cannot be evaluated only by whether it is technically useful. It must be evaluated by how it is governed. Public-interest infrastructure is not defined by nonprofit status, open-source code, public funding, or social-purpose language alone. Those features may help, but the more important questions concern authority, access, accountability, interoperability, stewardship, and protection against capture or enclosure. Infrastructure that influences how issues are understood before resources move carries power. That power needs governance.

The same principle applies to community-originated information. If community signals are essential to understanding an issue, then they must be handled as governance-bound knowledge rather than raw data. Community experience, local interpretation, cultural knowledge, trust signals, and lived reality cannot simply be collected, summarized, and converted into institutional evidence without attention to authority, consent, provenance, interpretation, protection, and benefit. A philanthropic institution may steward community-originated information under agreed conditions. It should not presume ownership of it.

The conclusion of this paper is therefore both practical and restrained. Philanthropy does not need infrastructure that promises perfect visibility, objective allocation, or certainty. It needs infrastructure that improves preparation. It needs a way to see more clearly before selecting organizations, designing grants, or moving capital. It needs a way to understand what is known, what is claimed, what is contested, what is missing, and whose authority shaped the interpretation. It needs a way to place organizations, grants, outcomes,

community signals, and capital flows inside the broader issue landscape rather than examining each in isolation.

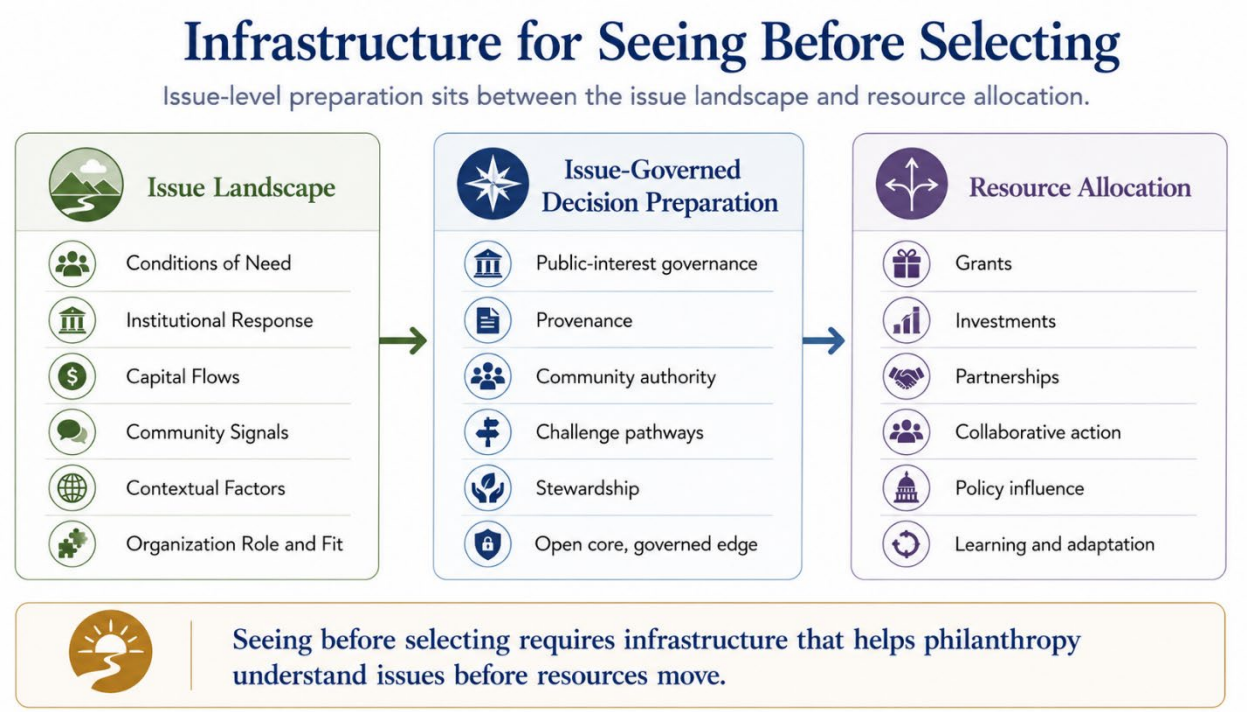

The deeper purpose of this infrastructure is not efficiency alone. It is accountability to the issue before resources move. An issue landscape exists whether funders can see it or not. Communities experience conditions whether those conditions appear in applications or not. Capital flows through systems whether those flows are mapped or not. Organizations operate within ecosystems whether those ecosystems are understood or not. The role of public-interest infrastructure is to make these realities more visible, more governable, more challengeable, and more responsibly connected to decisions.

This paper should be read as a companion to the first. The first paper asked what happens when philanthropy begins at the wrong level of analysis. This paper asks what would need to exist for philanthropy to begin at the right level more often and more responsibly. The answer is not another isolated tool. It is a governed preparation layer that helps the sector assemble issue-level knowledge before decisions are made.

Seeing before selecting is not only a better sequence. It is a different standard of responsibility. It asks philanthropy to understand the issue landscape before choosing organizations, to preserve the authority of community knowledge before using it, to make uncertainty visible before claiming confidence, and to build infrastructure that can serve the public interest beyond the priorities of any single funder, vendor, or institution.

That is the work before resources move.