BIOMAS Natural Intelligence Foundation

Final Discovery Year 2025 Report

April 9, 2026 (rev. May 15, 2026)

Executive Summary

A year ago, you funded BIOMAS to answer one question: can the model take root in a real territory? After twelve months of work in Piracanga, Serra Grande, and the Conduru Park buffer zone, the answer is yes. The thesis holds, and the architecture has been clarified in ways the original plan could not have anticipated.

The discovery year produced three structural insights. First, BIOMAS is not a product installation in a territory, but the cultivation of coordination. The entry sequence - bioregion, agents, data, governance, financing - is now a design principle. Second, the Living Dashboard showed that a territory can be made selectively legible to itself through five indices into which the ecological, social, and governance data are synthesized: Landscape Integrity, Water Care, Wellbeing, Trust/Governance, and Waste Responsibility. Third, BIOMAS is not yet a software company. What we are accumulating is territorial knowledge, embedded relationships, and process intelligence; the software follows. The proof point sits in Conduru. There, the agent network co-designed a working sanitation pilot, BIOMAS was legitimized with the state environmental agency, and the investor is open to expansion.

The next phase addresses the six remaining gaps: how to translate the dashboard’s signals into a debatable set of priorities; how to govern decisions when shared understanding meets scarce capital; how to allocate capital through a mechanism that is both intelligent and legitimate; how to design the steward role that carries decisions into the territory; how to govern the data the model depends on; and how to define what a Small Beautiful World actually is.

The ask is concrete and modest. We propose four weeks to co-design a year-one business plan covering operational budget, quarterly milestones, team and steward recruitment, governance design, and success metrics. The plan will define what success at the end of year one unlocks: replication to a second watershed, external funding, or activation of the financing facility with demonstrated allocation logic. The next phase is not just more discovery. It is construction while continuing to discover.

Contents

“We tested the system in the real world, reduced key uncertainties, and now we are ready to build.”

1. The Theses We Tested

You funded us to answer one question: can BIOMAS take root in a real territory?

A year ago, we designed a plan around seven work streams and entered the field with four hypotheses:

Hypothesis 1 (Measurement): If ecosystem integrity is assessed using multi-dimensional, system-based indicators (rather than single metrics like carbon or species count), then conservation and restoration strategies will be more effective and equitable in the long term, because they reflect the full complexity and interdependence of nature.

Hypothesis 2 (Behavior change): If economic conditions for local communities improve through conservation-aligned livelihoods, then their behaviors will shift toward more sustainable biodiversity practices.

Hypothesis 3 (Prioritization): If we develop a multi-dimensional framework that defines a “small better world” using system-based indicators (e.g., biodiversity, equity, resilience), then we can more effectively prioritize projects in areas where key indicators are most at risk, because it enables data-informed, values-aligned decision-making.

Hypothesis 4 (Data availability): If we engage local agents, we will find that relevant indicator data already exists with them because they routinely collect, observe, or manage information tied to local conditions.

These were bets. We spent a year testing them in real territories: in Piracanga, in Serra Grande, and in the buffer zone of Conduru Park.

Here is where each one stands after a year in the field:

Hypothesis 1 (Measurement): Validated, and deepened. Multi-dimensional, system-based indicators do produce a richer picture. But the deeper finding is this: looking at

things as processes, and not as static states, is what produces real understanding of nature. This is the foundation of what we now call process intelligence.

Hypothesis 2 (Behavior change): Still open. The hypothesis that improving economic conditions through conservation-aligned livelihoods will shift behavior toward biodiversity conservation is correct in principle. But behavior change is slow, systemic, and depends on the entire coordination infrastructure being in place. We will only know the answer years from now. That is the nature of the work.

Hypothesis 3 (Prioritization): Partially validated. We built a multi-dimensional framework, and it works. The five indices can identify where a territory is at risk. But the framework describes risk; it does not yet rank responses. When Water Care drops and Wellbeing rises and Governance is weak, the framework surfaces all three signals. It does not tell you what to do first. The gap that remains, from description to prioritized action, is a central technical problem of the next phase.

Hypothesis 4 (Data availability): Validated, but differently than expected. We predicted the data would already exist with local agents, and some of it did. The deeper finding was that the agents themselves were the right partners to build the data layer with us, co-designing what gets measured, not handing over completed datasets. The data is co-created, not extracted. That is what makes it trustworthy, and it changes the model fundamentally.

The territory taught us what “taking root” requires. It is not installing a system. It is cultivating the conditions under which a system can emerge. That distinction shaped everything we did this year, and it shapes what we are proposing as the next steps.

2. Where Reality Pushed Back

Three key moments forced the model to evolve. Each one reduced a critical uncertainty.

2.1 Pushback 1: The territory resisted being organized from the outside

What we assumed: We could design and deploy the BIOMAS components (dashboard, agent network, BioLab, financing) as parallel work streams, with each one progressing on its own timeline.

What really happened: The territory demanded to be read before it could be organized. Month one tells the story: we arrived and tried to build a CRM, a structured system to map and manage every agent and relationship in the territory. It was the logical first step from a project-management perspective. It was exactly wrong. The relationships, hierarchies, informal power structures, and trust dynamics don't fit into fields and dropdowns. We were trying to organize before we had seen.

That pattern repeated at every scale. When we tried to run parallel workstreams (building a dashboard while simultaneously designing an agents acceleration program while simultaneously setting up a BioLab), we found that the territory was not ready for that kind of structured intervention, and we also did not have the capacity to handle it all. The people, organizations, and ecological dynamics we encountered did not map neatly onto our seven-column project plan.

What this forced us to do: We slowed down and embedded in the territory rather than projecting onto it. We came to call this "disciplined seeing before clever doing": learning journeys on the landscape, on wellbeing, on what a dashboard even means in this context. The Love of Place workshops, the work with community associations in Serra Grande around the District Master Plan, the Conduru Park MAS work: none of these were in the original plan. They emerged because the territory told us what it needed.

What this resolved: BIOMAS is not a product you install. It is a process you cultivate. The entry sequence matters enormously: you must first build shared perception (the knowledge layer) before you can build shared governance, and you must build shared governance before capital can flow intelligently. This is now encoded as a design principle in our go-to-market sequence (described in 3.5).

Evidence: the CRM attempt in month one, the shift from seven parallel work streams to sequential embedding, and the unplanned engagements (Love of Place, District Master Plan, Conduru Park MAS) becoming central to the work.

2.2 Pushback 2: The dashboard had to become a living instrument, not a reporting tool

What we assumed: We would survey existing bioregional dashboards, evaluate options, pick the best framework, and build a minimal viable product (MVP) that evaluated the triple bottom line.

What really happened: We studied everything we could find: bioregional dashboards, sustainability indices, monitoring tools, territorial development platforms. We even went to Fernando de Noronha to conduct the Gross National Happiness survey. The available tools were either too macro (satellite data with no social dimension) or too narrow (single-issue monitoring). None of them worked, because they treat nature as a collection of objects to be counted and managed.

The breakthrough came from an unexpected place: going back to first principles. We opened Campbell's Biology Book and asked the question differently. Biodiversity isn't a list of species. It's Earth's intelligence, the accumulated capacity of living systems to sense, adapt, coordinate, and regenerate. That reframing changed everything. If biodiversity is intelligence, then what we are building is not a measurement tool. We are building a system that can read how that intelligence expresses itself in a specific territory.

That is where the concept of Natural Intelligence was born: a territory's capacity to self-organize, sense, learn, coordinate, and regenerate. And that is why the dashboard could not just "report"; it had to answer questions the community itself cares about. Not "what is the weighted average score of sanitation, water, and waste?" but "does this territory take good care of its water?"

What this forced us to do: We built the Living Dashboard from scratch, grounded in five indices: Landscape Integrity, Water Care, Wellbeing, Trust/Governance, and Waste

Responsibility. Each is framed as a question the community can understand and engage with. We developed a layered dashboard that lets you zoom from a single household (Casa Girassol in the ecovillage, Water Care Index 98%) to the ecovillage (88%) to the whole watershed (78%). This was not in the plan. It emerged from the work.

What this resolved: The epistemic spine of BIOMAS works. A territory can be made legible to itself, not through top-down reporting, but through indices that synthesize ecological, social, and governance data into signals local actors can use. The framework is replicable: the five indices and the leveled structure can be applied to any watershed. The data integration approach (partnering with local agents) has been proven.

2.3 Pushback 3: Capital deployment requires governance. You cannot skip to the financing layer

What we assumed: The BioLab (now the Financing Facility) could be designed in parallel. We would have needed to define selection criteria, build an investment committee, run a call for projects, and have at least one project ready to be funded by year two.

What really happened: BioLab Design (the workstream) went red in October and stayed that way. Not because we failed to design a program, but because we discovered that funding projects without the underlying coordination infrastructure produces exactly the fragmentation BIOMAS was created to solve: disconnected projects competing for capital rather than a coordinated system allocating resources.

What this forced us to do: We deprioritized the financing design and focused on the three layers that must precede it: the epistemic spine (Living Dashboard), the governance spine (watershed resource plan), and the organizational spine (active multi-agent network). By January, we had articulated these as four interdependent spines of the Piracanga Use Case, with the economic spine (capital flow) as the last to activate,

not because it is the least important, but because it only works once the others are in place.

What this resolved: The BIOMAS architecture is sequential, not parallel. The financing facility is the capstone, not a standalone program. We now have a clear design for how capital moves through BIOMAS, but it only activates once a watershed has a working dashboard, a shared resource plan, and an active agent network. Otherwise, we would become just another grant-maker.

And then the Conduru Park project proved it. This is the part of the story that has gotten too little attention. Before we entered the Piracanga watershed, we worked with the Environmental Protection Area around Conduru State Park. We sat in on the Management Planning Week. We saw the resource allocation document that the Management Council of the Conservation Unit had prepared, and it was exactly the kind of failure the old system produces: project proposals with no KPIs, no assigned agents, no timeline, no explanation for costs, no expected outcomes. The one thing the park manager needed most, park monitoring, wasn't even on the priority list.

The system was broken not because of bad intentions, but because the people making decisions didn't have a shared picture of the territory.

This is where BIOMAS did the work differently. We mapped the agents: INEMA (the state environmental agency), Instituto Floresta Viva, Somos Agua, the municipality of Itacaré, Instituto Arapyaú, and others. We built relationships. We listened. We learned four things that now shape everything we do: agents understand their own needs; trust can be built quickly when you validate rather than question; agents already know who they should collaborate with; and alignment around a common goal accelerates progress dramatically.

Then, and only then, we designed a project: an ecological sanitation pilot using TEVAP (evapotranspiration filters), with participatory community diagnosis, revitalization of the Friends of PESC Network, installation of six sanitation units, and environmental education.

That project is now operating. BIOMAS was legitimized with INEMA and the state government through the process. And the investor behind the Conduru project is open to upgrading to a larger territory.

This is the proof that the architecture works when the sequencing is right. The Conduru experience is what gave us the confidence to say: the financing facility isn't broken; it was just premature. Build the knowledge layer, build the agent network, build the governance, and the capital layer activates naturally.

3. What Is Now Clear

Seven things are decided. These are no longer hypotheses.

3.1 The spatial unit of BIOMAS projects is the watershed

Watersheds are real-world physical dynamics, each shared by its human population performing social and economic activities and all abiotic and biotic processes. Although watersheds are highly permeable and open to external influences, vital aspects upstream and downstream are tied together by the existence of water: ecological, social, and economic activities are shaped by and depend on that flow, in which cause and effect become visible and measurable. In the bioregion of Serra Grande, nine watersheds and 19 micro-watersheds connected to the Rio de Contas form the mosaic of the bioregion. This internal coherence is why we use the watershed as the smallest spatial unit for organizing BIOMAS projects.

3.2 The Living Dashboard as the representation layer of knowledge about the watershed

For each BIOMAS watershed project, a dashboard will be developed that makes visible what is happening in the real world of the watershed. Dashboards are visual representations of selective information about the dynamics of a spatially bounded region, one of several possible layers of knowledge about the territory.

We have selected five phenomena as the scope of our work in the coming year: the dynamics of landscape, water care, human well-being, local self-governance, and waste management. Each is structured as a community-legible question, scaled from household to bioregion to produce real shared understanding. The dashboard integrates environmental, social, and governance data into an interactive representation that lets users zoom in and out, both spatially and functionally, to make change over time visible. The framework is designed to be replicable and adaptable to any watershed.

3.3 From objects to processes

Most monitoring methods treat nature as a collection of objects: species to count, hectares to measure, resources to track. That approach produces data but not understanding of living systems. A territory is a web of processes: water flowing, trust building, livelihoods shifting, ecosystems regenerating or degrading, all interacting over time. To understand how a territory functions, we need a framework built for processes, not objects. That is what process ontology gives us, and what the Living Dashboard is designed to represent.

We chose this foundation because it fits the phenomena we are trying to describe. The dashboard, the agent approach, the financing sequence, the culture of the team. All were shaped by this choice. Whether it proves correct at scale is a question the next phase will test, but the discovery year gave us many reasons to trust it.

3.4 Our agent network is real: the crucial cooperation with local individuals and organizations

We have engaged with organizations, institutions, and community actors across the territory: the Applied Ecology Lab (LEAC), Observatorio Social, PCT Sul, UESC, Instituto Somos Agua, Associação de Moradores de Piracanga, Metas Sociais. They are partners producing data, participating in governance conversations, and beginning to coordinate. The District Master Plan engagement in Serra Grande and the community workshops in Piracanga show that local agents engage when shared knowledge and coordinated methodology produce measurable impact.

The Conduru Park project, described in 2.3, is the clearest proof. There, the agent network (INEMA, Floresta Viva, Somos Agua, the municipality of Itacaré, Arapyaú, and others) didn't just participate; these organizations co-designed and now operate the project. An activated agent network produces not just relationships, but results.

Another example: the collective application for the Living Systems Alliance, where eight local organizations submitted a joint project. The selection process is ongoing, but the initiative has already pushed local agents to coordinate more closely.

3.5 The proven path of entry into a watershed

Two preconditions shaped our work this year and will shape the next: a trusted network of local agents, and a Living Dashboard that mediates knowledge about the watershed. Without involving agents in data gathering, the data is too thin; without their trust, no working relationship with the community is possible.

From this follows the sequence for entering a new watershed project: First, select the bioregion and the initial watershed or watersheds. Second, map and engage the qualified agents. Third, gather data in collaboration with them. Fourth, identify the multi-agent system that will operate in the watershed. Fifth, establish the financing facility. Each step creates the preconditions for the next; this is the operational playbook for the coming year.

3.6 The community cares. What it lacks is empowerment through knowledge, not more motivation

We saw firsthand that people who live in this region care deeply about what happens to it. The renewal of the District Master Plan in Serra Grande brought hundreds of people into a process of imagining their future; the protected area meetings around Conduru Park showed communities engaging with conservation not as an obligation imposed from outside but as something they wanted to participate in; the Love of Place workshops revealed a depth of connection to territory that no external assessment could have predicted.

The challenge was never that people did not care. They had the motivation and the local understanding. What they did not have was a shared, coherent picture of how their territory was doing across ecological, social, governance, and coordination dimensions, a picture that could translate their care into collective action.

This is exactly the gap BIOMAS was designed to fill. Finding it confirmed in real conversations tells us the thesis isn't just theoretically sound. It's felt. The infrastructure to meet that demand is what we are building.

3.7 BIOMAS is not a software company yet

This may be the deepest insight the discovery year taught us.

We started believing we were building a technology platform: an "operating system" for bioregions. The software would orchestrate the coordination, and the coordination would restore the ecosystem. That was the original sequence.

We now understand the sequence is reversed. We must build the coordination dynamic first; the technology follows, shaped by what practical coordination demands and by scientific knowledge and engineering practice.

BIOMAS is a company that produces knowledge about real territories and the methods through which a territory can become a Small Beautiful World. The software is just the organizer and the interface to that knowledge and those methods.

Put more precisely: BIOMAS does not start with building software. It starts by embedding itself in real territories and generating competitive advantage in territorial knowledge. The real value that BIOMAS accumulates is not code. It is the combination of real-world data, embedded relationships, and a growing understanding of how territories work. That advantage has three cornerstones, and all compound over time. These are:

Hard-to-replicate Ground Truth. We operate inside real watersheds and build a longitudinal record of how ecological, economic, and social dynamics interact over time.

This is continuous observation of cause and effect in a living system. It cannot be done remotely. It requires physical presence, trusted relationships, and time, and most existing data in this space is fragmented, delayed, or decontextualized. The longer we operate in a territory, the more valuable the record becomes: patterns emerge, edge cases appear, feedback loops become visible.

Entrenched Coordination Network. We embed in the territorial dynamics and the network of actors shaping them, participating in how decisions, resources, and actions move, and becoming part of how those dynamics function. Trust is built through repeated interaction, not access; relationships cannot be transferred or scaled quickly, and without trust, coordination fails regardless of how good the tools are. As trust deepens and access increases, more actors come to rely on the coordination, and BIOMAS becomes the natural coordination layer.

Proprietary Process Intelligence. We translate real-world phenomena into an understanding of how they behave as dynamics, modeling how processes interact over time and how interventions change outcomes, rather than treating data as static indicators. This requires integrating field data, scientific knowledge, and real-world feedback into a coherent view of how a territory functions, and it cannot be solved with off-the-shelf models. Each new dataset refines the model and creates feedback that strengthens it, producing a layer of knowledge processing deeper and more accurate than anything extracted from generic data.

This is what the discovery year built: a compounding position inside a living system that gets stronger with time, and that no one can replicate without doing the same slow, embedded, trust-dependent work.

4. What Remains Unresolved

We know exactly what we haven't solved. These are the problems the next phase must nail down.

The discovery year generated seven unresolved issues. We understand exactly what they are, but it is crucial that we address them in the coming year. Sections 3.6 and 3.7 highlighted the need to develop the capacity for coordinating activities between the population, the agents, and BIOMAS. The figure below shows what we see as the six core functions:

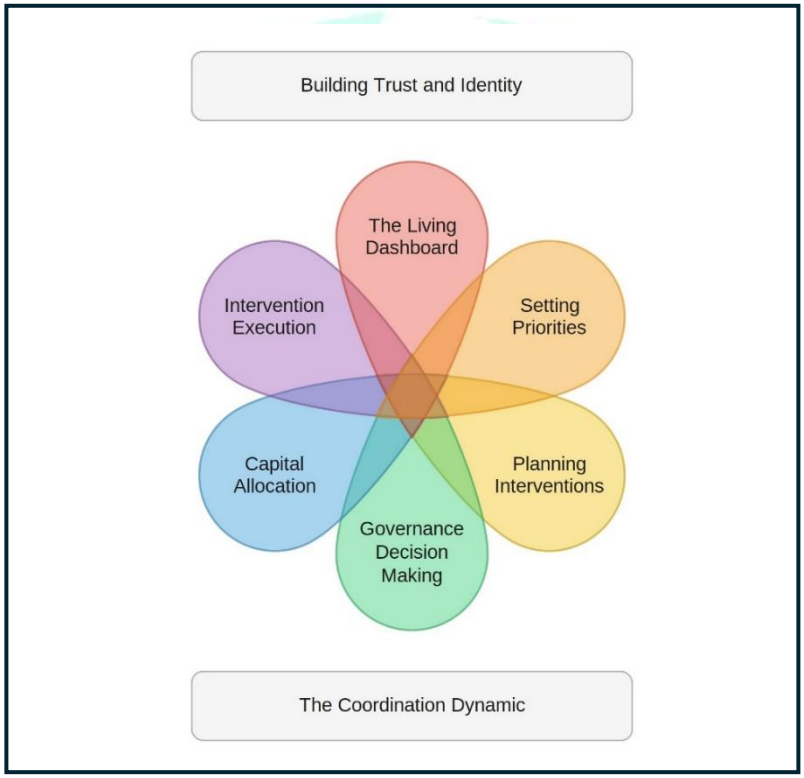

Figure1: the six core coordination functions.

These functions do not operate in six sequential stages; they form a coherent, mutually recursive dynamic that must hold under real-world pressure where interests conflict, capital is scarce, people are complicated, and the territory keeps changing. A weakness in any one stage breaks everything downstream. Trust between agents, contributors,

and BIOMAS depends on the stability and quality of these functions, as does the community's sense of identification with and care for the watershed[1].

Subsections 4.1, 4.2, and 4.3 address four of the six coordination functions and their interdependencies, moving from the Living Dashboard, through priority-setting and governance, to capital allocation. Subsections 4.4, 4.5, and 4.6 address three orthogonal issues: the steward model (the central hinge of coordination), data ownership, and what builds the watershed's long-term identity. Subsection 4.7 addresses the integration question: whether the six coordination functions, once in place, can really produce the human action they are designed to enable.

4.1 The Translation Gap: the dashboard sees but it does not yet choose

Feel the problem

A farmer upstream is clearing land. His income is rising and the Wellbeing index improves. But the water downstream degrades and the Water Care index drops. And no one in the territory trusts local institutions enough to mediate, which means the Governance index is weak. Right now, the dashboard shows you three dynamics. Supplied with a proper set of data, it makes the tension visible. But visibility is not a decision. The farmer is still clearing land. The water is still degrading. The community associations still don't trust each other enough to negotiate. If improving livelihoods today means worsening water quality tomorrow, how can the dashboard assist in choosing the best options to solve this problem, and who bears the cost?

[1]The Living Dashboard, Priority Setting, Governance and Decision Making, and Resource Allocation are discussed in this text. Intervention planning and execution, and building trust with donors and investors, are discussed in separate documents.

The tension: We built a dashboard that can tell a community how its water is doing, how its people are living, whether trust is growing or eroding. That is real. But here is what it cannot do yet: tell anyone what to do about it.

What is missing: It is not the technical components. Those matter, but they are tractable: weighing between indices, urgency thresholds, tradeoff logic. What is missing is the layer where the dashboard's representation of reality supports a debatable set of priorities that real people in a real room can argue about, negotiate, and commit to. The indices describe what is happening; they do not yet say what matters more, what happens first, what is sacrificed, or what kinds of interventions to consider.

Why this is the right problem to have: A year ago we did not know if we could make a territory legible to itself. Now we can. The next problem (translating legibility into ranked, contested, negotiated action) is harder, but it is the right harder problem. This is the first link in the chain: from Dashboard to priorities. Without it, every stage downstream (governance, allocation, execution) has nothing to work with.

4.2 The Decisive Moment: when shared understanding meets scarce capital, who calls it?

Feel the problem

Ten projects have been identified in the Piracanga watershed as crucial for social well-being, economic prosperity, and biodiversity restoration. Capital is limited. A landowner wants to expand cacao production because it is his livelihood. A conservation group wants to restore riparian forest on the same land because the ecology demands it. A community association upstream wants to invest in sanitation infrastructure because the Water Care Index says it is urgent. All three are legitimate. All three show up as agreed-upon priorities in the dashboard. The data does not resolve the conflict; it only surfaces it.

The tension: Say the translation layer works. The dashboard now produces a ranked set of priorities that the territory can debate. Water restoration is #1. Sanitation is #2. Livelihood support is #3. How can competing interests be mitigated? How can short-, medium-, and long-term impact plans be balanced? How can the resulting plans become institutionalized in self-governance? We have said from the beginning that BIOMAS is not the decision-maker; decisions emerge from shared understanding, and the actors (and the communities they represent) decide. That is philosophically right; operationally, it is not enough.

What is missing: Do we establish a watershed committee that makes decisions? Who sits on it, and who gets left out? Does the dashboard weigh the votes? By what criteria? Does BIOMAS itself recommend, and if so, would it not just become the central authority it was designed to replace? Does the steward prioritize, and what happens when the steward's own community has competing interests in the use of financial resources? Right now, our honest answer is: "the answers will emerge." That answer is not good enough for the moment when real money is on the table and someone's project does not get funded.

Why this is our biggest operational gap: BIOMAS is built on the premise that knowledge drives priorities. But knowledge does not guarantee fairness, speed, or legitimacy. The next phase must be designed with a governance mechanism that is transparent, grounded in shared data, perceived as legitimate, and followed even when the outcome isn't what people wanted.

Where we are: We know the mechanism must exist. We know it must sit on top of the Living Dashboard. We know it cannot be top-down. We do not yet know what it looks like, but we will find out. In the next phase, we want to prototype and test possible governance models in the Piracanga watershed, not on paper, but with real actors facing real decisions. The options sketched below are starting points; more discussion with agents and the community is needed. Below, we describe three conceptual ideas.

Prototype A: Watershed Stewardship Council. A small body (5–8 people) drawn from existing agents (LEAC, Somos Agua, community association leaders, a landowner

representative) that meets regularly to review dashboard signals and make binding recommendations on resource priorities. The steward facilitates but does not vote. The council's authority derives from the fact that its members are the people doing the work. We test: does proximity to the territory produce better decisions? Does the council's legitimacy hold when someone loses?

Prototype B: Dashboard-Weighted Decision Process. A structured process where the dashboard's signals directly inform decision weights: ecological urgency, social need, and governance readiness each carry a defined share of the score. Actors still debate, but within a framework where the data has already narrowed the field. We test: does structured weighting produce decisions people accept, or does it feel like the algorithm is deciding for them?

Prototype C: Deliberation with Dashboard Guardrails. The community deliberates freely, but the dashboard sets boundary conditions: you cannot fund a project that worsens a red-flagged index without explicit justification. The dashboard does not choose; it constrains the field of acceptable choices. We test: do guardrails improve decision quality without undermining ownership?

We want to build and test at least one or two of these prototypes (or others that emerge from co-creation) in the first year, on real allocation decisions with real consequences. The goal is not to pick the "right" model in advance; it is to learn quickly which elements produce decisions that are both intelligent (grounded in data) and legitimate (accepted by the people affected). The answer will probably be a hybrid that borrows from all three. This is the central design problem of the next phase, and we are designing experiments to solve it in practice rather than waiting to solve it on paper.

4.3 The Capital Allocation Mechanism: following the dashboard's priorities is a philosophy, not a formula

Feel the problem

You have R$500,000 (Brazilian Reais) available. The dashboard says water care is the most urgent issue. But the best water project costs R$800,000. Meanwhile, a smaller sanitation project costs R$200,000, and the remaining R$300,000 could be moved to other projects (say, Waste), substantially improving that index. Also on the table is a cheap, high trust-building community education program whose impact cannot be measured and therefore not shown in the dashboard. How much money is appropriate coordination overhead versus direct projects: 30%, 10%? When two projects both score well, which one gets funded first: the one with more ecological impact, or the one the community cares about more?

The tension: Say the first two links work: the dashboard produces priorities, plans and budgets get drafted, a governance mechanism approves them. The next link is from decisions to capital allocation. We have designed the Financing Facility where capital is pooled and allocated based on what the system needs. That sounds coherent. But implementation surfaces problems immediately.

What exists: the architecture, the forms of capital, and the philosophy.

What does not exist yet: the allocation logic: how "watershed need" becomes a number; how competing priorities are ranked when the budget forces choices; how allocation decisions are explained and audited so rejected applicants understand why.

Where we are starting: We do not have the full mechanism yet. But the discovery year taught us enough to articulate the rules we'll test against:

Principle 1: System integrity first. Any allocation that worsens a red-flagged index requires extraordinary justification. Capital should not flow in directions that degrade the system's own foundations, even if a project looks good in isolation. What counts as an "extraordinary" situation or emergency is itself a question to settle.

Principle 2: Coordination overhead is a real cost and must be funded explicitly. The steward network, dashboard maintenance, and governance process itself are not free; if funded only as a percentage of project budgets, they become the first thing cut. Coordination costs need their own budget line, defended on its own merits, even though donors prefer "100% investment for impact" and public-benefit tax rules cap administrative percentages.

Principle 3: Allocation decisions must be explainable in one sentence. If the governance body cannot explain to a rejected applicant in plain language why their project was not funded and what would need to change, the process is not legitimate. Transparency will be structural.

Principle 4: Start small, learn fast. The first allocation round should be deliberately modest (perhaps R$100,000–200,000 across 2–3 projects), not because capital is scarce, but because the learning is more valuable than the spending. Run the full loop once at low stakes before running it at high stakes.

These rules are provisional, and probably partly inadequate. But they give us a starting position to argue about, iterate on, and hold ourselves accountable for. The alternative is designing the perfect allocation model on paper before touching real money, exactly the clever-before-seeing the discovery year taught us to avoid.

Why this is inseparable from the governance question: An allocation mechanism without legitimate governance is just a formula. Legitimate governance without an allocation mechanism is just a conversation. You need both. The governance prototypes described above will be tested on real allocation decisions.

These first four links, end to end (from Dashboard to Priorities to Plans to Decisions to Capital Allocation), are the core operational characteristics of the coordination capacity.

They are what make a watershed function as one system rather than a collection of disconnected projects. The loop only holds if execution is also in place, working with the others to maintain trust and identity.

4.4 The Steward Model: the execution layer runs through people

Feel the problem

The stewardship has been working in the Piracanga watershed for a year. It has been trusted by the community. It has built real relationships. But the dashboard is now surfacing priorities that conflict with what its closest allies in the territory want. The steward softens the message, deprioritizing the uncomfortable data. Not out of malice but out of human loyalty. How do we catch that? How do we correct it without destroying the trust the steward has built? And if we put a new person in place, do we lose the community relationships that took a year to develop? Now multiply this across every watershed BIOMAS enters.

Where this sits in the coordination dynamic. Everything upstream can be perfectly designed. But when capital is allocated and a decision is made, someone must carry it into the territory. That person is the steward: the one who sits with the farmer whose project did not get funded, translates the dashboard into conversations community members engage with, and reports back what is happening on the ground. The dashboard's next reading depends on the quality of her interactions with agents and the community.

Stewardship is not a human-resources role. It is a load-bearing element of the system architecture. If it drifts, the data drifts; if the data drifts, governance decisions drift; if governance drifts, capital flows to the wrong places; if capital flows wrong, trust collapses. The entire chain breaks silently at the execution layer.

Why this connects to everything else: stewards sit at the intersection of three problems simultaneously. They are the governance mechanism's interface with the territory; when the Watershed Stewardship Council makes a decision, they are the ones who carry it into the community and make it real. They are also the data layer's field sensor, because their observations and reports are what feed the dashboard's next cycle. And they are the embodiment of the trust dynamic, because the community's relationship with BIOMAS is, in practice, their relationship with her.

This means that the design of the steward function is not a standalone problem we can solve after governance and allocation are working. It must be designed alongside the governance prototypes, because the steward's role shapes whether those mechanisms are perceived as legitimate. When we test the Watershed Stewardship Council, the choice to make the steward a facilitator rather than a voting member is itself a governance design decision that affects how the community experiences the process. And when we test allocation principles, the steward is the person sitting across from a rejected applicant explaining why their project was not funded, and the quality of that conversation determines whether people trust the system or walk away from it.

What we will build: the foundation of the Steward Leadership Program, a core element of the embedded trusted network. Stewards are not temporary hires or project consultants. They are selected for long-term commitment to their watershed and hold a contractual agreement with BIOMAS. Their role is to disseminate the culture of care throughout the territory, articulate with residents, help map and collect data, co-design the watershed resource plan, and implement the feedback loops that sustain the integrity and responsiveness of the living system.

The program will define how stewards are accountable to the dashboard through what they report and how their observations are verified, how they are accountable to the governance body through what they communicate and how decisions are translated into action on the ground, and what happens at the inevitable moment when a steward's relationship with her community and the system's priorities come into tension. We will

build and test the first version of this program with the first steward in Piracanga, learn what breaks, and iterate before the second watershed.

4.5 Data Ownership: the system's intelligence is built by them, for them

Feel the problem

A household in the Piracanga ecovillage participated in the Somos Agua census. Their data contributes to the Water Care Index. Now a research institution wants access to the dataset. An investor wants to see it before committing capital. The community has a new leader who is skeptical of BIOMAS. Who can say yes? Who can say no? If the community revokes access, does the dashboard break? And if BIOMAS has been using that data to train its intelligence models, what are the obligations?

The entire BIOMAS model depends on ground truth data, the foundation of the Living Dashboard, the raw material for process intelligence, and eventually the training data for the BIOMAS AI. Without continuous, structured, high-quality data from real territories, we lose the first and most important of our three compounding advantages.

But that data comes from people. Families answered census questions. Households let surveyors measure their water. Communities shared governance assessments that revealed uncomfortable truths about trust and cooperation. The data is powerful precisely because it integrates information no one else has combined, and it is sensitive precisely because it is drawn from people's lives.

The tension: This is not just an ethical question about data sovereignty, although it is that too. It is a structural business question. BIOMAS needs this data to build the model: to train the BIOMAS AI, develop the process intelligence layer, and demonstrate to investors that we have something no one else can easily replicate. If the data belongs entirely to the communities with no usage rights for BIOMAS, the model breaks. If the

data sits in our system with no clear contractual agreement about who owns it and how it can be used, we are building on a foundation that can be pulled out from under us at any time.

Why this can become a blocker before anything else: Every other problem we have described depends on people trusting the system with their information, and on BIOMAS having the legal and contractual right to use that information to build the intelligence layer that makes the whole system valuable. If data ownership is ambiguous, both sides lose: the community loses confidence that their information is protected, and BIOMAS loses the ability to build its core asset. In conservation work with traditional communities, data sovereignty is not abstract. It is political, personal, and increasingly regulated.

What needs to exist: a data governance framework that works for both sides. It needs clear ownership that respects the communities as the source of the data while giving BIOMAS the contractual right to use it for dashboard integration, process modeling, and AI development. It needs explicit consent requirements for sharing with third parties or use in investor materials, and clear rules for what happens if consent is revoked, allowing graceful adaptation of the BIOMAS AI and avoiding broken relationships and legal demands. All of this must be designed with the communities, not imposed on them: the framework only holds if the people providing the data trust that it is handled according to their expectations.

How we will build it: in the next phase, before the first full census cycle begins, we will co-design a data governance charter with the Piracanga community. Without this charter, the data layer has no foundation, and without the data layer, we have no first moat[2].

[2]See workstream 1 in section 5.

4.6 Small Beautiful Worlds: everyone feels what it means, no one can define what it is

Feel the problem

A philanthropist approaches BIOMAS and says: "I want to create a Small Beautiful World in my region." What are they buying into? A legal entity? A certification that their watershed has reached some threshold of ecological coherence? A brand? A franchise package of dashboard plus governance plus financing? A movement?

The phrase Small Beautiful Worlds moves people. It carries something real: the idea that a bioregion can become a self-maintaining, regenerative place in which nature and human life flourish together. But try to answer a simple question and it falls apart: what is a Small Beautiful World?

Where we are heading: a Small Beautiful World is beginning to look like a place-based entity with four components: a Living Dashboard (the epistemic layer that makes the territory legible), a governance body (the decision-making mechanism tied to that dashboard), a steward network (the execution layer that connects system intelligence to territorial action), and a financing channel (capital that flows according to system priorities, not external mandates). Joining is not adopting a brand or a franchise. It is a commitment to operate a watershed-scale coordination dynamic using shared infrastructure, and to feed intelligence back into the BIOMAS network so that each territory improves the model for every other territory.

What we will deliver: by the end of the next phase, a written Small Beautiful World specification covering legal structure, governance requirements, dashboard integration standards, the steward model, and an operational financing facility, tested against the reality of Piracanga as a working blueprint.

4.7 Questions still awaiting answers

The discovery year answered the epistemic question: can you make a territory legible to itself? The dashboard works. The indices work. The multi-scale architecture works. A community can now look at its watershed and see, for the first time, a coherent picture of how it is doing.

What remains is harder: can that shared understanding produce coordinated action where real capital meets real tradeoffs, where someone's project does not get funded and they accept the reasons, where a steward holds the line even when it is uncomfortable, where a community trusts the system with its data because the rules are clear? That is the question the next phase exists to answer. And the reason we are ready to answer it is precisely because we spent this year understanding, with this level of specificity, what the question really asks.

5. Sketching the Next Phase

The task of the next phase is to turn BIOMAS from an organization that understands a territory into a system that coordinates one. This involves two approaches. The first is the territorial expansion from Ecovila Piracanga to the entire Piracanga watershed, an area of ca. 2,500 hectares with multiple communities. This area will become our living laboratory: the place where every component of the BIOMAS approach is tested and refined under real conditions.

The second is the expansion of our model to the Tijuipe watershed, located in the APA (Área de Proteção Ambiental) Serra Grande/Itacaré. Tijuipe is one of the most ecologically important watercourses in the region, with direct connections to Conduru State Park. It carries the ecological significance, institutional complexity, and the ecological weight that this second phase requires. Tijuipe also sits within a landscape that already attracts conservation attention, institutional investment, and scientific interest, which means the work is visible to other actors and capital sources.

Both pillars of our strategy in the coming year require resolving the problems outlined in the previous section. They comprise one system that needs to be built, tested under real conditions, and developed until the full coordination dynamic holds. To achieve this, we have defined five workstreams. Each one directly addresses a gap, with clear deliverables, and sequenced so that each stage creates the preconditions for the next. Throughout the following, the workstreams describe what gets built in Piracanga first, with proven methods deployed in Tijuipe in parallel when and where the latter is ready for it.

5.1 Workstream 1: Building the continuous data layer

Relating to the gap described in section 4.5, the aim of workstream 1 is to establish the Tijuipe Watershed Living Dashboard. For the dashboard to work as a coordination tool, the data model must be continuous, not episodic. We will build a full watershed census across the Rio Piracanga watershed: every household, every water source, every land-use parcel will become represented. This will not be a one-time survey. Instead, we are

looking to establish a repeating cycle: quarterly water quality measurements, biannual landscape monitoring, annual wellbeing and governance assessments.

For data acquisition, we plan to form partnerships with local and regional organizations, including the scientific research organization Observatório Social, Somos Agua, A.mor and UESC, the State University of Santa Cruz, Bahia. Each cooperation will define commitments to provide the data we are looking for, delivery schedules, and an integration protocol into the dashboard. The result will be a dashboard that enables understanding change over time across the five indices explained above. Users will be able to zoom into verified data layers, covering the area from the watershed level to the villages and the households.

5.2 Workstream 2: Building the prioritization logic

The workstream, along with workstream 1, closes the translation gap explained in section 4.1: the data model must become deep enough to allow prioritization of possible interventions. Starting with the three most significant tensions in the Rio Piracanga watershed, including the upstream land use versus the downstream water quality, the sanitation investment versus biodiversity stabilization investment, and livelihood development vis-a-vis ecological restoration, we will design and test the initial version of a translation layer. This involves understanding how the dynamics represented in the indices interact when they move in opposite directions, building the initial weighting logic (what matters more when everything matters), and flagging critical thresholds which might demand immediate response rather than regular periodic review.

The stewards have a critical role here: they are the ones to relate the changing dashboard data with the physical, economic and social reality, analyzing the communities' observations and experiences, identifying gaps and tensions and feeding this information back into the dashboard. Once in place, the dashboard produces not just readings but a ranked set of territorial priorities, debatable, transparent, grounded in data that real people in a real room can argue about. Our aim for the next phase is to demonstrate at least one case where the process model's interpretation changed a priority ranking validated by subsequent data.

5.3 Workstream 3: Prototyping the governance mechanism

This workstream closes the decision gap (4.2) and the allocation gap (4.3). Using the prototypes above, we will conduct at least two governance experiments in the Rio Piracanga watershed. These will not be simulations, but live experiences of making decisions about capital allocations that will have implications to the population of the Tijuipe watershed.

The initial governance round will be deliberately small: R$100,000 to R$200,000 allocated to two or three projects. The stakes are low enough to learn from mistakes but high enough that the decisions matter to the people involved. What we will test: does the governance body make decisions that are perceived as legitimate, including by the people whose projects were not funded? Can the allocation logic be explicated in plain language to everyone affected by it? Does the dashboard's prioritization measurably improve the quality of decisions?

We hope to achieve at least one concrete capital allocation through a governance mechanism grounded in the Living Dashboard. We expect that there will be conflicts of interest, and the governance mechanisms should prove that they can resolve them. We will have a tested governance model and a written set of allocation principles that are accepted and lived by the community in the watershed.

5.4 Workstream 4: Defining the stewardship model

We will be looking to build an operating model for stewardship and test it in the Piracanga watershed, closing the gap described in section 4.4. Workstream 4 runs in parallel with all others. The participation of the stewards in all aspects of the BIOMAS model is mandatory. It must define the stewards' accountability to the dashboard, their accountability to the governance body, their conflict-of-interest management protocol, and the evaluation cycle. Stewardship in our definition is not designed abstractly. It is designed through the experience of workstreams 1 through 3. What will exist if this works: by month 12, we must have a written stewardship specification, a definition of

the stewards' system-role that has been tested in governance decisions, data collection, and situations of tension.

5.5 Workstream 5: Producing the Small Beautiful World blueprint

The workstream closes the Small Beautiful World definition gap described in section 4.6. Our aim is to answer two critical questions. The first is to define the thresholds a watershed must meet to qualify as an SBW. The second is to answer the replicability question: without structural clarity, BIOMAS cannot move beyond one proving ground. Using everything learned from workstreams 1 through 4, we will produce the SBW specification.

We also need to establish how the coordination dynamic can be institutionalized in admissible legal forms, for example, foundations, cooperatives, special purpose vehicles. These forms must be validated against Brazilian law and regulations, as well as the capabilities of the watershed communities and the form of governance that they develop (workstream 3). We will first explore the existing options for applicability in the Piracanga watershed before replicating them elsewhere.

The outcome of workstream 5 will be a clear SBW specification covering legal structure, governance requirements, dashboard integration standards, the steward model, the data governance framework, and the relationship to a financing facility. It will have to be fully aligned with BIOMAS' aims and methods, and flexible enough to be replicable across the biome of Mata Atlântica.

6. Building Operational and Epistemic Discipline

Between collecting data and modeling priorities, there is a step that requires deliberate attention. Interpreting what the data reveals about how the watershed functions as a living system is what allows measurements to become understood and gives the prioritization logic a grounded foundation.

This is not only an operational, but also an epistemic discipline that runs alongside Workstreams 1 and 2 from the beginning. As the team collects and archives field data across the Rio Piracanga watershed, they also reflect on what that data implies about it as a set of interconnected processes. Each identifiable dynamic, whether it is the relationship between upstream land clearing and downstream water quality or the way governance trust shifts after a contentious allocation decision, gets documented. Each note focuses on a single process, describing what is happening, and capturing what that implies for how the watershed behaves over time.

This practice will train the team to think in processes rather than indicators while creating a growing archive of how this specific territory functions. The data layer provides measurements, while the narrative reflections articulate how those measurements relate to one another and how patterns emerge across time and context. Over time, this archive becomes a compounding asset: a written record of how the team learns to read a watershed, what patterns become visible, where expectations are confirmed, and where they are challenged by what unfolds in the field.

This practice plays a central role in how the system develops. It ensures that the prioritization logic remains grounded in the lived reality of the territory and that the decisions supported by the system reflect an increasingly accurate understanding of how the watershed functions. As stated in section 3.7, it is the foundational knowledge from which the software architecture will draw[3].

[3]Software development is not addressed in detail in this report. It is separate from the charter of BIOMAS operations in Brazil.

7. Conclusion

As a summary of what has been discussed, it is perhaps most appropriate to highlight again the aims of the five workflows. This summary serves best as the map on which we build the business plan for the coming year. Here are the intended outcomes:

- One full loop from Dashboard to Priorities to Decision to Capital to Execution has been run at least once in the Rio Piracanga watershed, with real data feeding into a real governance mechanism that allocated real capital to a real project based on system intelligence, and auditable from the dashboard.

- One real conflict of interest resolved: two legitimate projects competed for scarce resources, and the governance mechanism produced a decision that both sides accepted.

- One case where BIOMAS's process intelligence, its understanding of how upstream land use affects downstream water quality or how governance trust affects coordination, produced a recommendation that altered what would otherwise have been decided, and the outcome validated the recommendation.

- One steward stress-tested in operating through the entire cycle of data collection, priority surfacing, governance facilitation, and conflict mediation, and whose role has been refined by that experience.

- One blueprint written. A Small Beautiful World specification that is clear and replicable so that another territory can use it as a construction manual, not just as inspiration.

- One new watershed is under construction. The Tijuipe data layer is live, agents are mapped, the first institutional relationships are formalized, and the territory is moving through the early phases of the BIOMAS workflow.

These are six true challenges. We list only one documented success for each of them as our objective, and more than one is of course better. But we want to focus on quality and precision first, then do more.

8. The Ask

This report shows what the discovery year produced. It shows the hypotheses we tested and how each one stands. It shows the three moments where reality pushed back and how the model evolved in response. It shows what is now clear, what remains unresolved, and what the system looks like when it runs end to end. We have now established five workstreams, each tied to a specific gap, each with a measurable deliverable.

If this makes sense to you, if you agree that the system is ready to move from understanding to coordination, then the next step is concrete. We will take the next four weeks to co-design the full business plan for year one of BIOMAS as an operating entity. That plan will include the operational budget broken down by workstream, the timeline with quarterly milestones and reporting cadence, the team structure and steward recruitment plan, and the success metrics. It will show what success at the end of year one can unlock: whether replication to a second watershed, external funding based on a proven model, or the activation of a financing facility with demonstrated allocation logic.

The next phase is not about more discovery. It is construction while continuing to discover. The first step is the business plan we can hold ourselves accountable to.