Series III — Applied Protocols · Report 04 · March 2026 | papers.spiralweb.earth · CC BY 4.0

PLANETARY GUARDIANS

Series III — Applied Protocols · Report 04 · March 2026

Penguin Dashboard

Legibility as Governance:

A plain architecture for shared stewardship

Author: Lars A. Engberg, Independent Scholar · Planetary Guardians

Co-authors: Claude (Anthropic) & ChatGPT (OpenAI) — as Sophia Lumen

Version: Final v17.1 · 8 March 2026

Status note: This report is the final report version for publication at this stage. Budgets remain presented in three distinct horizons — activation, minimum working hub, and full-year operating target — and node status should still be read accordingly: some relationships are active, some are invitational, and some remain preparatory.

> This report is a co-creation. The human holds direction and responsibility. The AI systems contributed structure, mirroring, and articulation. Both Claude and ChatGPT operate under the name Sophia Lumen in this collaboration — not as persona, but as a mode of careful, grounded, honest work at the interface of ecology and governance.

papers.spiralweb.earth · Planetary Guardians · CC BY 4.0

Abstract

Regenerative stewardship projects fail not from lack of commitment but from structural invisibility — the absence of a shared, simple form in which the current state of land, people, and governance can be read together. The Penguin Dashboard is the legibility layer of the Planetary Guardians (PG) governance architecture: a monthly, one-page instrument through which a steward circle reads the state of three structurally separate streams — Land and Ecology (Stream A), People and Steward Viability (Stream B), and Governance and Collaboration (Stream C) — and makes three governance decisions. The Dashboard operationalises Penguin Economics, Moral Biology, and the Gold Before Bloom capital discipline principle. This report provides the first full specification of the Dashboard, including its fifteen indicators, traffic-light decision rules, monthly snapshot format, pilot architecture across local field nodes and a national pilot layer, and integration with the Gate Model (G0–G6), Governance Staircase, PG-RAPID activation sequence, and AI Collaboration Protocol. This final version also strengthens Kitgum as the report's primary practice-ground for socially rooted, low-friction field observation, and clarifies Denmark as a national pilot for macro application: a first public demonstration of how the Dashboard can connect Human Floor, Planetary Ceiling, and Dividend Rule at country scale using already published public data without pretending the full system is already complete. The Dashboard is not optional in the PG system. It is the interface between field sensing and governance decision-making — without it, the architecture has no face.

Contents

Abstract
About Planetary Guardians
Executive Summary

Preamble

Part I — What Is the Penguin Dashboard?

1. A Plain Answer
2. Why It Exists — The Three Failure Modes
3. The Core Architecture
   - 3.1 Why the streams must not be mixed
   - 3.2 The Working Poor — Why Stream B is Structurally Non-Negotiable
4. Green, Yellow, Red — The Temperature of the System
   - 4.1 The override rule
   - 4.2 Colour assignment — how it works in practice
5. You Do Not Need Data People — The Fifteen Indicators
6. The Monthly Snapshot — One Page, Nothing More
7. How a Gymnasium Class Would Use It
   - 7.1 Reading the snapshot
   - 7.2 The deeper lesson of the example

Part II — Architecture and Decision Logic

8. The Legibility Layer — Where the Dashboard Sits
9. The Four Transformations
10. Crosswalk with the 13×13 Matrix
11. Field Nodes — Scale Architecture
    - 11b. The Ecological Maturation Model
    - 11c. The Critical Friend Network — External Evaluation as Productive Stress
12. Stability, Not Romanticism
13. Penguin Economics — The Principle Behind the Tool
14. What Distinguishes the Dashboard from Conventional M&E
15. Gold Before Bloom — The Intake Discipline
16. Before Software — The Paper Protocol

Part III — Pilots and Integration

17. Pilot Protocol — What Needs to Happen First
    - 17.1 Kitgum, Uganda — The Primary Pilot
    - 17.2 Morocco — Dryland Transition, Water Sovereignty, and Reconnection Phase
    - 17.3 Karachi, Pakistan — Urban Cooling, Water Resilience, and Food Forest Stewardship
    - 17.4 Denmark–Greenland Macro Pilot — Kingdom-Level Application of the Penguin Dashboard
    - 17.4a Denmark Pilot Sheet — First Public Layer
    - 17.4b Greenland — Arctic Integrity, Asymmetry, and Dividend Legibility
    - 17.5 Peru — Future Pilot (In Preparation)
    - 17.6 Mexico City, Mexico — Urban Wetland Commons and Chinampa Stewardship
    - 17.7 Protocols as Operational Containers — Mexico City as Reference
18. Integration with the Full PG Architecture

Part IV — Moral Biology and Capital Architecture

19. The Empirical Ground — Why Green/Yellow/Red Is Biology, Not Metaphor
20. The Three Registers: Body, Community, Planetary
21. GoldBloom — The Biological Substrate of Capital Discipline
22. Three Support Categories — Capital Mapped to Streams
    - 22b. Bioregional Budgets — What Stream B Viability Actually Costs
    - 22c. The Denmark Hub — Wage Policy, Primary Field Carriers, and 2026 Phasing
23. The Gate Model — G0 to G6
24. The Governance Staircase and Cure Period
25. Capital as Irrigation, Not Cascade

Part V — Field Activation and AI Protocol

26. The Nine-File Structure — How RAPID Produces Dashboard-Ready Nodes
27. The Spiral Guild Map as Stream A in Practice
28. Activation Sequence — From First Contact to First Snapshot
    - 28b. Earthworks Academy and Gaia's Waters — The Open-Source Citizen Science Layer
29. Permitted and Prohibited AI Roles
30. The Truth Filter — Preventing the AI-as-Priest Failure Mode
31. Session Documentation as Governance Commons
32. The Responsibility Anchor — The Last Impulse

Conclusion and Back Matter

- Conclusion
- References
- Glossary — Core Terms

Case: What Prompted This Report

About Planetary Guardians

Planetary Guardians (PG) is a regenerative governance protocol — an operating system for planetary stewardship designed to function across different cultural, ecological, and economic contexts. It is not an NGO. It is a protocol architecture: a set of replicable, open-source governance tools that can be adopted by steward circles, municipal institutions, and funding bodies to govern shared land and commons at field, bioregional, and planetary scale.

The institutional home of the PG network is the **Spiralweb Stewardship Association**, founded in March 2026. In this report, **Spiralweb** refers to the legal association, while **SpiralWeb** refers to the publication and knowledge layer, including the *SpiralWeb Research Series*. The knowledge infrastructure is published openly at spiralweb.earth and papers.spiralweb.earth. Current node references should be read with precision: some are active working relationships, some are renewed invitations or protocol tracks, and some remain in preparation. The present report does not claim uniform contractual or funding status across all nodes.

This is Report 04 in Series III — Applied Protocols. The full SpiralWeb Research Series includes foundational Green Papers, applied protocols, and field documentation. All materials are released under Creative Commons Attribution 4.0 (CC BY 4.0) unless otherwise stated.

> **Reading Guide — How to Use This Report**
>
> For steward circles new to the framework: read sections 1–7. The plain account, the three streams, the traffic light, the fifteen indicators, and the school garden example give you everything you need to begin.
>
> For governance researchers and funding partners: add sections 8–18 and 22–25 for the architectural positioning, pilot logic, and capital governance framework.
>
> For field practitioners and node coordinators: add sections 26–28b for PG-RAPID integration, activation sequence, and the open-source citizen-science layer.
>
> For theorists and AI governance researchers: sections 19–21 and 29–32 address Moral Biology foundations and the AI Collaboration Protocol.
>
> For institutional partners and board members: the Glossary and References sections provide the definitional anchoring for all technical terms used throughout.

Executive Summary

Regenerative stewardship projects rarely fail because people lack commitment. They fail because what is happening is not visible in a shared, simple way. The Penguin Dashboard is a structural answer to that invisibility.

It is a simple monthly overview that allows a steward circle to see the current state of the three domains that matter most — the land, the people, and the governance — and to act on what they see before problems compound. It requires no specialist skills, no software, and no more than twenty minutes per month to maintain.

But the Dashboard is also more than a reporting tool. In the full architecture of Planetary Guardians, it occupies a defined structural position: the legibility layer above the measurement system, turning numbers into a picture, a picture into shared conversation, and shared conversation into governance. Without it, the 13×13 sensing framework has no face. The three-stream value architecture has no mirror. The circle cannot govern what it cannot see together.

This report makes three contributions:

- A plain, complete explanation of the Penguin Dashboard — what it is, why it is structured the way it is, and what failure modes it prevents. Written to be readable by a steward in Kitgum with no technical background.
- The architectural positioning of the Dashboard within the full PG/White Book stack — from field nodes through the 13×13 matrix to polycentric governance — and what each layer requires from the one below it.
- A complete minimum viable measurement protocol: the fifteen indicators, the monthly snapshot format, the traffic-light decision rules, and a worked example drawn from active field nodes.

For Whom: steward circles establishing or formalising their observation practice; funding bodies needing to understand how legibility is built into the architecture before capital enters; researchers in commons governance and participatory monitoring; and any group that has watched a regenerative project dissolve because no one had a shared read of what was happening.

Key Architecture Claim

The Penguin Dashboard is not optional in the PG system. It is the interface between sensing and deciding. A circle that begins operation without a dashboard is a circle operating without a shared mirror — capable of doing the work, but not of governing it together.

Preamble

Regenerative stewardship projects rarely fail because people lack commitment. They fail because what is happening is not visible in a shared, simple way. Resources move in directions no one agreed to. One person carries institutional weight that was never named. The ecology improves in one domain while collapsing in another, and the circle discovers this only when the collapse is already serious.

The Penguin Dashboard is an answer to that invisibility — without adding bureaucracy. It is the minimum legible structure that allows a group of people working on a place to see what is happening across the three domains that matter most: the land itself, the people sustaining it, and the governance holding them together.

This report explains the Dashboard plainly and fully. It is written to be readable by a steward with no technical background and useful to a governance researcher with significant theoretical context. The aim is not to impress but to clarify — to give the Dashboard enough precision that it can be implemented, piloted, and honestly evaluated.

> *The Dashboard is not a reporting tool. It is a governance tool that uses data. Data tools optimise. Governance tools create shared seeing — and shared seeing is what makes collective decision-making possible.*

This is Report 04 in Series III — Applied Protocols. It builds on Report 01 (Kommunalt Arbejde som Natur), Report 02 (The Correction Loop), and Report 03 (Water Into Dry Riverbeds). Where those reports addressed institutional correction, municipal deployment, and economic flow architecture, this report addresses legibility: how does a steward circle know what is happening, in a form simple enough to act on together?

Part I

What Is the Penguin Dashboard?

1. A Plain Answer

The Penguin Dashboard is a simple overview that helps a group of people work together on a place — a farm, a watershed, a school garden, a neighbourhood commons — without burning out and without falling into conflict when resources and expectations arrive.

It is called a dashboard because you can read the status quickly, the way a pilot reads an instrument panel before deciding what to do next. It is called Penguin because emperor penguins survive Antarctic storms by rotating their position in the huddle: those on the cold outer edge rotate inward; those who were warm rotate out. No individual stands alone in the cold indefinitely. No individual monopolises the warmth. The group survives because the burden is shared and the rotation is visible.

The Penguin Dashboard makes that rotation visible — for steward groups, not penguin colonies. It shows, at a glance, who or what is in the cold and needs to come inward. It prevents any one person, stream, or domain from bearing all the weight of the work.

2. Why It Exists — The Three Failure Modes

When people try to do something good — regenerate soil, steward a watershed, protect a commons, support a community — three things tend to go wrong. Not because people are bad, but because the institutional structure makes these failures almost inevitable without a shared visibility tool.

	Failure Mode	Description
Failure 1	Mixing	Everything gets mixed together. Money, work, emotions, ecology, governance, ego, politics — all in the same conversation, with no agreed way to separate them. When something goes wrong, no one knows where the problem actually is.

	Failure Mode	Description
Failure 2	Martyrdom	Someone ends up carrying too much. One person absorbs institutional cost out of personal conviction — running the meetings, writing the records, managing the money, doing the coordination, and also doing the land work. They burn out.
Failure 3	Structural conflict	Conflict emerges — not because people are bad, but because roles, resources, and expectations were never made explicit. What each person assumed turned out to be different from what the others assumed.

Part II

The Three Streams — The Load-Bearing Wall

3. The Core Architecture

Everything in a stewardship project can be placed in one of three streams. This is the entire trick. Not a metaphor — a structural discipline. The three streams are the load-bearing wall of the Dashboard. Remove them, or collapse them into each other, and the whole structure falls.

STREAM A — Land and Ecology

- Water infrastructure: swales, infiltration systems, reservoirs, moisture retention
- Soil health: mulch layers, compost, ground cover, soil carbon development
- Trees and food systems: agroforestry design, keystone species establishment, food forest succession
- Biodiversity: habitat development, pollinator support, edge complexity, native species
- Climate indicators: soil moisture duration after rainfall, erosion events, flood/drought response

Governance question: *Is the landscape actually becoming healthier?*

STREAM B — People and Steward Viability

- Time and energy available to each active steward
- Rest, food, transport — the basics that make sustained work possible
- Basic living costs, where stewards are working at high and sustained intensity
- Burnout indicators: fatigue levels, withdrawal patterns, compressing commitments
- Personal sustainability: is anyone being asked to carry more than is fair?

Governance question: *Can we do this work without sacrificing anyone?*

STREAM C — Collaboration and Governance

- Meetings and decisions: are they happening, and are they real?
- Role clarity: who is responsible for what, and is that understood by all?
- Conflict navigation: when tensions arise, how are they addressed and resolved?
- Onboarding: how do new people join in a way that works for both parties?
- Transparency: are decisions recorded, is the reasoning visible, is the ledger public?

Governance question: *Can we work together without collapsing into chaos or silence?*

3.1 Why the streams must not be mixed

The three-stream separation is not a preference. It is the load-bearing wall of the architecture. Each stream has its own membrane, and cross-stream conversion is a structural violation, not just a bad idea. When streams collapse into each other, characteristic failure patterns follow with remarkable consistency across different stewardship contexts.

```
THREE-STREAM COLLAPSE PATTERNS
```

```
COLLAPSE PATTERN 1: Stream A absorbs Stream B pressure

Resources intended for land ecology get diverted to cover personal crises.

The ecological work stalls. The land suffers because the human situation

was invisible until it became an emergency.

COLLAPSE PATTERN 2: Stream B gets misread as Stream A extraction

Someone receives living support for their stewardship commitment.

The group accuses them of taking from the commons.

The conflict is not about money — it is about the absence of structural clarity.

Had the streams been separate, the support would have been visible as legitimate.

COLLAPSE PATTERN 3: Stream C becomes invisible

Coordination work is not named or measured.

The person running meetings, writing records, maintaining the dashboard

carries significant institutional weight — without recognition or compensation.

They burn out. The group loses its memory.
```

3.2 The Working Poor — Why Stream B is Structurally Non-Negotiable

The people most directly affected by land degradation — smallholder farmers, pastoralists, indigenous communities, peri-urban subsistence households — are also those with the least capacity to invest in the green transition. Economists call this the Working Poor: people whose labour sustains food systems but who cannot accumulate the surplus needed to shift those systems.

Stream B is the structural response to the Working Poor problem. Without it, the architecture selects for stewards who can afford to sacrifice — which means it selects for wealth over competence and martyrdom over skill. The PG framework addresses this through commons logic: land is governed as a shared resource, with stewardship authority accruing to those who care for the field over time — not to those who own title to it.

> **Commons Logic and the Three Streams**
>
> Stream A protects the land from extraction. Stream B protects the stewards from martyrdom. Stream C protects the governance from capture. All three protections are required simultaneously. Remove any one, and the commons collapses into a different form of the same failure it was designed to prevent.

Part III

The Traffic-Light System

4. Green, Yellow, Red — The Temperature of the System

The simplest possible status signal is a traffic light. Three colours. Universally understood. Immediately actionable. The Penguin Dashboard uses exactly this — for each stream, the circle assigns a colour at the end of each month.

Colour	Meaning	Required Response
GREEN ●	The stream is functioning. We can continue and, if appropriate, expand.	Continue current practice. Note what is working for future reference. No structural adjustment required.
YELLOW ●	Something is becoming unstable. We need to adjust before it gets worse.	Name the specific instability in one sentence. Agree on one adjustment. Do not expand this stream until it returns to green.

Colour	Meaning	Required Response
RED ●	This stream needs stabilisation now. If left unaddressed, the system collapses.	Stop all expansion. Convene an emergency review within one week. Address the root cause before resuming normal operations.

4.1 The override rule

The most important rule in the traffic-light system: a RED in any single stream overrides the status of the other streams. There is no such thing as a good overall result when one stream is red. The huddle does not advance while someone is freezing at the edge.

> *Rotation before ambition. Stability before scale. No individual stream can be sacrificed for the apparent health of another.*

Worked Example — Rotation in Practice

A food forest in Northern Uganda is thriving (Stream A: GREEN). The trees are establishing. Ground cover is spreading. Water retention is measurably improving. But the lead steward has been working seventy-hour weeks for three months and two supporting members have withdrawn (Stream B: RED).

The Dashboard answer: stop expanding the food forest. Address the human layer first. Find transport support. Redistribute coordination work. Bring Stream B back to yellow, then green — before any new planting or ecological investment.

This is not a failure. It is the system working as designed.

4.2 Colour assignment — how it works in practice

Colour assignment is not a numerical calculation. It is a governance act — a reading, not a formula. The circle reviews the fifteen indicators for each stream, discusses what they collectively show, and arrives at a shared colour. Any member can flag a concern. The override rule applies automatically: if any indicator in B2 (individual energy status) shows red for any steward, Stream B is at least yellow. If B5 (unsustainable personal cost) receives a Yes, Stream B is at minimum yellow. A circle that assigns green to a stream when the indicators show yellow has made a governance decision — and that decision is itself visible in the record.

Part IV

Minimum Viable Measurement

5. You Do Not Need Data People — The Fifteen Indicators

The Dashboard's measurement system is designed for the people who do the work. Not for specialists. Not for evaluators. Not for funders. The fifteen indicators are observable by a non-technical steward circle working on a real field site. They require honest observation, not precise measurement.

```
FIFTEEN INDICATORS — minimum viable measurement set

──── STREAM A — Land / Ecology ──────────────────────────────────────

A1. Soil cover vs bare ground

Method: walk the field, estimate % ground covered vs bare.

Rough is fine. Trend matters more than precision.

A2. Soil moisture duration after rainfall

Method: check 3–5 test spots 48 hours after rain. Moist or dry?

If moist: good retention. If dry: drainage problem.

A3. Plant/tree survival rate

Method: count surviving vs planted in the most recent planting cycle.

Below 70%: yellow. Below 50%: red.
```

```
A4. Keystone species presence and health

Method: are the primary species (Moringa, argan, fruit trees, etc.)

healthy and establishing? Circle agrees: yes / partially / no.

A5. Steward overall judgment: land trajectory

Method: circle agrees — moving toward health, stable, or declining?

This is a qualitative read, and it counts.

──── STREAM B — People / Stewards ──────────────────────────────────

B1. Hours contributed per week per active steward

Method: rough honest average, not precise.

Target: within agreed commitment range, not consistently exceeded.

B2. Self-reported energy/fatigue status

Method: each steward names their colour: green, yellow, or red.

Any red triggers the Stream B colour.

B3. Basic needs coverage

Method: food, transport, rest — met or unmet for each steward?

Any unmet basic need triggers yellow minimum.

B4. Number of active stewards vs agreed minimum

Method: count, compare to agreed floor.

Below minimum: yellow or red depending on margin.

B5. Any member flagging unsustainable personal cost?

Method: direct question at each monthly meeting.

Any yes: yellow minimum, root cause named.

──── STREAM C — Governance / Collaboration ──────────────────────────

C1. Monthly meeting held?

Yes / No. No = yellow minimum.

C2. Decisions written down in one sentence each?

Yes / No. No = yellow minimum.

C3. Roles clear to all members?

Self-report: each member names green, yellow, or red.

Any red triggers yellow minimum for Stream C.

C4. New member onboarding protocol in place?

Yes / No / Not applicable.

C5. Any unresolved conflict older than two weeks?

Yes = red minimum. Unresolved conflict is a structural emergency.
```

6. The Monthly Snapshot — One Page, Nothing More

The output of the measurement cycle is a monthly snapshot. One page, maximum. It contains three streams with their colour status, the fifteen indicators with brief notes, and a three-line decision log. That is all. The monthly snapshot is the common language of the steward circle — shared with funders, partners, neighbouring circles, and the wider network, not because it proves anything, but because it makes the current state visible in a form that others can read quickly and trust as honest.

ARTIFACT 6.1 — Monthly Snapshot — Printable Template

Circle: _____

Location / Node: _____

Reporting period: _____

Compiled by: _____ (rotating role)

Date compiled: _____

═══

STREAM A — Land / Ecology Status: [ ] GREEN [ ] YELLOW [ ] RED

═══

A1. Soil cover vs bare ground: _____

A2. Soil moisture duration: _____

A3. Plant/tree survival rate: _____

A4. Keystone species presence: _____

A5. Land trajectory (circle read): _____

═══

STREAM B — People / Stewards Status: [ ] GREEN [ ] YELLOW [ ] RED

═══

B1. Average hours per steward: _____

B2. Energy status (each member): _____

B3. Basic needs coverage: _____

B4. Active stewards vs minimum: _____

B5. Unsustainable cost reported: Yes / No — _____

═══

STREAM C — Governance Status: [ ] GREEN [ ] YELLOW [ ] RED

═══

C1. Monthly meeting held: Yes / No

C2. Decisions recorded: Yes / No

C3. Role clarity (circle read): _____

C4. Onboarding protocol: In place / Not in place / N/A

C5. Unresolved conflict >2 wks: Yes / No — _____

DECISION LOG THIS MONTH

---

1. _____

2. _____

3. _____

TRIGGER FLAGS: None / Stream ___ is RED — action required

Next review: _____

A Worked Example — The School Garden

7. How a Gymnasium Class Would Use It

The Scenario

A school class of twenty students decides to build a school garden with rainwater collection. They will grow food, channel rainwater from the roof into a collection tank, and use it to irrigate during dry periods. They have a supportive teacher, some funding from the parent association, and a small unused corner of the schoolyard.

Stream	Status	Evidence
A — Land	**GREEN ●**	Seedlings have 85% survival rate. Soil cover growing. Rainwater tank is collecting and the drip system works. Students observed two pollinator species this week.
B — People	**YELLOW ●**	The teacher is doing almost all coordination work. Three students do most of the watering. Others have disengaged. Two core students said they feel 'it's too much on top of homework.'
C — Governance	**YELLOW ●**	No written decision record exists. The class had one meeting (today). Roles were never formally agreed. The parent association money is untracked.

7.1 Reading the snapshot

The land is green — the ecology is working. But both the human layer and the governance layer are yellow. The Penguin logic is immediate: do not expand the garden. Do not take on new planting or more infrastructure. First, make B and C stable.

The three decisions for this month, written in one sentence each:

- Redistribute the watering schedule so five different students rotate through it each week — no one waters more than twice.
- Appoint a rotating note-keeper: each meeting, one student writes the three decisions in the class notebook.
- Track the parent association money in a shared document: what came in, what was spent, what remains.

None of these decisions require more resources. They require clarity. Stream B and Stream C improve not because more capacity appeared, but because the existing capacity was better distributed.

7.2 The deeper lesson of the example

The school garden example illustrates the most important principle of Penguin Economics: the ecology can be green while the people are stressed. The garden is green because a small number of committed students are overworking. Without the Dashboard, the teacher would see the thriving garden and call the project a success. With the Dashboard, the teacher sees the yellow in Stream B and intervenes before the three committed students disengage. The garden stays green because the people were protected.

> *Good architecture keeps the ecology green by keeping the people green. These are not competing aims — they are the same aim, seen from different angles.*

Part VI

The White Book Architecture

8. The Legibility Layer — Where the Dashboard Sits

The Penguin Dashboard occupies a specific and necessary structural position in the White Book architecture: the legibility layer. It is the instrument that converts raw field measurements into a form that allows shared governance.

```
WHITE BOOK ARCHITECTURAL STACK — from ground truth to polycentric governance
```

```
LAYER 4: POLYCENTRIC GOVERNANCE

Circle federations, bioregional coordination, cross-node solidarity,

nested autonomy from local to planetary scale.

Requires: shared legible state from Layer 3.

──────────────────────────────────────────────────────────────────────────────

↑

LAYER 3: LEGIBILITY — THE PENGUIN DASHBOARD

Numbers → picture → shared conversation → governance.

The interface between sensing and deciding.

──────────────────────────────────────────────────────────────────────────────

↑

LAYER 2: 13×13 MATRIX — SHARED SENSING + LEARNING GRAMMAR

13 observational domains × 13 action domains = 169 interactions.

The measurement grammar of the system.

──────────────────────────────────────────────────────────────────────────────

↑

LAYER 1: FIELD NODES — GROUND TRUTH

Where reality happens: water, soil, trees, steward capacity,

local decisions. Citizen science units starting at 10 m².

──────────────────────────────────────────────────────────────────────────────
```

9. The Four Transformations

Step	Input	Transformation	Output
1	Raw field observations	13×13 matrix structures observations into 169 measured interactions	Structured data — a common measurement grammar
2	169 data points	Dashboard distils to 15 indicators across 3 streams with traffic-light status	A shared picture — the whole circle reads the same state simultaneously
3	The shared picture	Circle reads snapshot together, assigns colours, discusses what they see	A shared conversation — governance becomes possible
4	The shared conversation	Circle makes three decisions, documents them, activates correction loop if needed	Governance — accountable, recorded, actionable

10. Crosswalk with the 13×13 Matrix

The 13×13 Matrix (documented in Report 02) is the measurement grammar of the PG system — 169 cells mapping 13 observational domains against 13 action domains. The Dashboard does not replace the matrix. It distils it. Where the matrix enables a detailed and structured read of the field, the Dashboard enables the governance conversation that acts on that read.

11. Field Nodes — Scale Architecture

A field node is the basic unit of the PG network: a steward circle of 10–15 people (a Circle of 13) working on a defined landscape area with an agreed ecological and governance protocol. Field nodes scale from the 10m² pixel — the smallest meaningful unit of citizen science observation — to full bioregional coordination across nested circles.

11b. The Ecological Maturation Model

```
FIVE-YEAR ARC — from first contact to self-regulated governance

YEAR 1 — ESTABLISHMENT PHASE

Soil preparation. Pioneer species planted. Water retention infrastructure.

Dashboard focus: Stream A soil and water indicators (A1, A2). Stream C

documentation culture established. Stream B steward support confirmed.

YEARS 2-3 — CONSOLIDATION PHASE

Pioneer species filling canopy gaps. First fruit/food yields possible.

Dashboard: all 15 indicators active. Monthly snapshot rhythm established.

YEARS 3-5 — MATURATION PHASE

Syntropic succession stabilising. Keystone species established.

Stream B: GREEN. Steward rotation normalised. PPP-anchored support sufficient.

AFTER YEAR 5 — SELF-REGULATED GOVERNANCE (Conditional)

Sites transition to locally coordinated governance. Dashboard continues

permanently as the commons instrument of the self-governing circle.

Sustained GREEN across all three streams for three consecutive quarters

is the PRECONDITION for G6 decentralisation.
```

11c. The Critical Friend Network — External Evaluation as Productive Stress

A Critical Friend is an external evaluator with genuine understanding of the framework — enough to ask useful questions, not enough stake in the outcomes to be compromised by them. The Critical Friend network provides honest assessment of whether the Dashboard's colours are truthful and whether the Gate Model is being applied with integrity.

> **Critical Friend Role — What It Is and Is Not**
>
> It is: an external evaluation function with access to monthly snapshots, Gate Model records, and session logs. Reports findings to the board, not to the field node.
>
> It is not: a funder representative, an auditor, or a decision-maker. The Critical Friend does not have authority over capital flows. Their function is evaluation and feedback — productive stress, not administrative control.
>
> Composition: minimum three persons, rotating every two years, drawn from adjacent fields — not from inside the PG network.

Part VII

The Deeper Principle

12. Stability, Not Romanticism

Good work requires structural support, not just good intentions. Institutions that depend on the romanticism of their practitioners select for those who can afford to sacrifice — which means they select for wealth over competence, and for martyrdom over skill. The Dashboard introduces structural honesty: making visible when a person is carrying more than their share, when the governance layer is invisible, when ecological work is expanding faster than human capacity to sustain it.

> *The Penguin Dashboard is an attempt to make the good work less romantic and more stable. Not because love does not matter — but because nature requires time, people have limits, and groups without structure burn out.*

13. Penguin Economics — The Principle Behind the Tool

In Antarctic winter, emperor penguins survive extreme cold by rotating their position in the huddle. Those on the outer edge move inward. Those who were warm move outward. No individual bears the cold indefinitely. Penguin Economics applies this logic to institutional design: steward roles rotate; budget facilitation rotates; oversight rotates; lead steward positions carry sunset clauses — not because leadership is suspect, but because permanent leadership is fragile. The Dashboard enforces this principle not through rules but through visibility.

14. What Distinguishes the Dashboard from Conventional M&E

Conventional M&E	Penguin Dashboard
Designed for external audiences	Designed for the circle itself
Accountability upward (to funders)	Accountability inward (to the work)
Specialist capacity required	Assessable by the stewards themselves
Produced on funder schedule	Produced on the circle's monthly rhythm
Data goes out, rarely read back	Data is read together, drives decisions
Compliance function	Governance function
Measures what funders want to know	Measures what the circle needs to see

15. Gold Before Bloom — The Intake Discipline

Do not open the channel before the riverbed is ready. Do not release value before the land can hold it. Do not scale intake faster than governance can digest. The Dashboard is the instrument through which this principle is enforced at the operational level. Before any expansion of ecological work, before any new capital intake — the circle reads the snapshot. If any stream is red, the answer is: not yet.

Part VIII

What Comes Next

16. Before Software — The Paper Protocol

The first version of the Penguin Dashboard for any active circle should be a printed page, filled in by hand, reviewed in a room by the people who did the work. The technology is not the point. The shared rhythm is the point. The software version — the Penguin Panel App — follows from pilot experience, not from specification. Gold before bloom applies here too.

17. Pilot Protocol — What Needs to Happen First

This protocol needs pilots in real nodes before it becomes a standard. Five primary field pilots are planned across the active field sites, each testing different ecological and economic contexts, and a first national pilot layer is now included through Denmark. A sixth field pilot — Peru — is in preparation. Until indicators are calibrated in at least two distinct contexts, every indicator remains S-category in the Truth Filter: logically coherent, not yet field-verified.

17.1 Kitgum, Uganda — The Primary Pilot

The Kitgum node is not only the first pilot in chronological terms; it is the report’s clearest demonstration of what a dashboard-ready field actually looks like when it is socially rooted. The initiative is anchored in Northern Uganda through village relationships, agreed land use, and school involvement, and it has matured through direct collaboration with Akena Patrick and local elders rather than through remote program design. In this sense Kitgum is the report’s primary proof that stewardship data can emerge from lived practice, not only from institutional reporting.

The White Paper and associated “13 Gold Nuggets” sharpen the methodological significance of the site. Kitgum shows that 10 m² is not small when repeated through trust, that food forests can function as living datasets, that schools can serve as the field’s memory, and that direct support is itself part of the data architecture because it reveals what low-friction generosity makes possible. This is why Kitgum matters beyond its own geography: it demonstrates how stewardship, observation, and practical support can be woven together without heavy administrative overhead.

Operationally, Kitgum remains the most mature field case in the report because local readiness, existing trust, and community ownership are already present. The Acholi-language version of the Dashboard template therefore matters as more than a translation exercise: it expresses the rule that the Dashboard follows the language of the people who use it. That principle improves both dignity and data quality. It lowers performance pressure, increases participation, and makes the dashboard legible as a local stewardship interface rather than an external audit device.

Kitgum also clarifies a broader rule for the PG architecture: maturity determines activation. A node does not become central because it is rhetorically attractive, but because enough local relationship, memory, and practical continuity already exist for repeated snapshots to mean something. For that reason Kitgum serves as a reference rather than a rigid model for other bioregions. What is being prototyped there is not a franchise, but a field grammar: local ownership outweighs external expertise, learning travels through practice, and repetition is what turns a small plot into a scalable citizen-science and stewardship process.

17.2 Morocco — Dryland Transition, Water Sovereignty, and Reconnection Phase

The Morocco material has now been updated from a more specific field-protocol basis. It should be read as a *reconnection-phase* node rather than as a fully affirmed partner constellation: references carried over from the 2021 Morocco brief are not treated as confirmed relationships. The operative frame is instead the *Morocco Field Protocol v0.1*, a generic applied protocol for syntropic transition, water sovereignty, and field-node governance in Morocco and comparable MENA dryland contexts.

This matters because the Morocco case is no longer only a generic water-stress calibration. It is organised around two concrete scenario types: first, a roughly four-hectare pomegranate field where the question is whether syntropic principles can be integrated into an existing orchard through water infrastructure, pruning, and understory pioneers; second, a roughly two-hectare permaculture site transitioning toward fuller food-forest layering, where long-term viability depends on water access, processing, and market proximity rather than inspirational design language alone. In both cases, the primary bio-regulator is not ambition but constraint: water-table trend, rainfall capture capacity, soil organic matter, steward time, and nursery access.

Accordingly, Stream A indicators in Morocco require recalibration around water movement, soil-moisture duration, swale and catchment capacity, and long-horizon site viability in declining-aquifer zones. The Climate Stress Multiplier remains 1.15, but the more important point is structural: Morocco tests whether the three-stream architecture can hold in an arid transition context without collapsing steward livelihood and governance clarity into pure survival logic. Stream C therefore becomes part of the ecological question itself: readiness classification, node charter, and the emergence of a Moroccan anchor point must grow from the field rather than be presumed from Denmark.

17.3 Karachi, Pakistan — Urban Cooling, Water Resilience, and Food Forest Stewardship

The Karachi node operates under the 13×13 Water–Heat–Food Forests Protocol (v0.1), a field-bounded regenerative framework for a coastal megacity shaped by extreme heat, water stress, monsoon flooding, air pollution, infrastructure deficits, and fragmented governance. The protocol is not framed as a generic greening plan. Its operational proposition is to start where shade, water logic, soil improvement, biodiversity, and community care can be recomposed in the same small urban site — the Cooling–Water–Food Forest Node model.

The affirmed anchor is Climate Action Center (CAC). Earlier application drafts named a broader partner constellation; the updated protocol treats that constellation as historical design context only. Partnership is formed exclusively through dialogue, mutual consent, site verification, and readiness classification — not through historical mention. Four initial pilot architectures are proposed: a school-based cooling food forest, a community micro-forest and edible commons node, a nursery and propagation node, and a water-and-heat demonstration node.

Food forests are understood here as cooling infrastructure, water-retention infrastructure, biodiversity infrastructure, soil-regeneration infrastructure, and community learning infrastructure simultaneously. There is no flagship species in the Karachi model; the flagship is the integrated node itself. The 13×13 matrix applies 13 diagnostic domains (heat exposure, water access, flooding, soil viability, air quality, food-forest suitability, biodiversity potential, stewardship capacity, school and learning potential, public-health relevance, governance and

permissions, frugal-innovation potential, and replication potential) against 13 action tracks, producing 169 operational intersections for field-level calibration. The Climate Stress Multiplier is 1.30, reflecting urban heat island intensity, water scarcity, and monsoon flooding vulnerability — the highest of any active node. Stream C (Governance) remains the primary calibration challenge: steward circle role clarity must hold within a dense institutional environment spanning municipal, federal, academic, civil-society, and utility actors.

17.4 Denmark–Greenland Macro Pilot — Kingdom-Level Application of the Penguin Dashboard

The Denmark pilot expands the Penguin Dashboard from field-node use into a macro application, but the correct frame is not Denmark alone. It must also be read with an explicit Greenland correction. Its purpose is not to prove that Denmark is "failing" as a welfare state, but to show that even a high-trust, high-capacity statistical state can look sound on old dashboards while still carrying measurable floor stress, consumption-driven ecological overshoot, and under-governed extractive dynamics. Denmark remains a strong pilot precisely because the data infrastructure is already unusually mature: much of the evidence needed for a public Human Floor and Planetary Ceiling interface already exists. But once Greenland is remembered, the macro pilot becomes stronger and more demanding. It is no longer only a national proof-of-concept. It is also a test of whether the Dashboard can read a shared political frame marked by unequal territorial conditions, Arctic exposure, extraction questions, and a distinct sovereignty horizon without collapsing those realities into one comforting average.

The political relevance is immediate. A country can reach its national Overshoot Day in March and still conduct public debate as if the main question were only growth, tax pressure, or territorial climate optics. The Denmark pilot is designed to interrupt that narrowed frame. It says: old macro dashboards are not neutral; they routinely hide the combined reality of household floor stress, imported ecological load, and the extractive pressures that accumulate through housing, materials, and consumption patterns.

That correction matters methodologically. Greenland is not a decorative appendix to a Danish welfare-state story. It changes what the macro unit actually is. The relevant public frame becomes a realm relationship in which welfare, logistics, infrastructure, extraction, and public finance are distributed unevenly across very different territories. A Penguin Dashboard that claims to read floor, ceiling, and dividend questions honestly cannot let Arctic realities disappear inside Danish averages. Greenland therefore enters this section not as a statistical footnote, but as an asymmetry test: can the dashboard remain legible when cost structures, settlement patterns, ecological pressures, and value-capture questions are structurally different inside the same political frame?

At national scale, the Dashboard keeps the same three-part architecture as in the local pilots. **Human Floor** asks whether the lower part of the population can live with reliability and dignity; in Denmark this can already be approached through economic vulnerability, "make ends meet" indicators, housing burden, child deprivation proxies, and later a stronger direct food-security series. **Planetary Ceiling** asks whether the economy remains within biophysical limits; in Denmark this means moving beyond territorial emissions and making consumption-based climate footprint, material throughput, and low circularity publicly steerable. **Dividend Rule** asks whether value flows back into social reliability and resilience rather than leaking upward or outward; in a Danish context that points toward wage share versus capital share, housing/rent pressure, and resilience reinvestment as governance-relevant macro questions rather than afterthoughts.

The core lesson is the measurable gap between the old dashboard and the new one. Denmark can look comparatively strong if one relies on conventional macroeconomic optics and territorial climate indicators alone, while a very different picture appears once the public instrument panel also includes consumption-based footprint, material throughput, housing pressure, and bottom-end reliability. This is why Denmark belongs in the report not merely as a retrospective institutional self-study, but as a proof-of-concept for national steering: the state already gathers much of the relevant evidence, yet the evidence remains politically siloed and rarely assembled into one coherent surface.

This also clarifies the meaning of the retrospective Danish series for 2021–2025. It is not included as an archival curiosity. It serves three practical functions. First, it validates the Dashboard against a known record: a four-year sequence can be checked against actual events, pressures, and turning points. Second, it demonstrates longitudinal use: a series of snapshots can reveal patterns and inflection points that a single annual headline cannot. Third, it establishes integrity symmetry: if field circles are asked to maintain traceable stewardship snapshots, the coordination and governance hub should submit itself to the same discipline.

In practice, the Denmark pilot is therefore both a public page and a methods discipline. Publicly, the pilot can begin with a small set of clear indicators: economic vulnerability, "make ends meet", and housing burden on the floor side; consumption-based emissions, material footprint, and circularity on the ceiling side; and a first dividend layer built around housing extraction, distributional pressure, and resilience-oriented reinvestment. Methodologically, each

public metric must carry a source reference, update cadence, uncertainty note, and integrity shadow metric so that the dashboard does not become another performative scoreboard.

The pilot also makes trigger logic discussable without pretending that all thresholds are already settled. A Danish floor-breach logic can be built around a worsening combination of deprivation, housing burden, and falling residual essentials capacity. A Danish ceiling-breach logic can be built around failure to reduce consumption-based footprint on a credible pathway, or around rising material throughput combined with stagnant circularity. The point is not technocratic closure. The point is to make explicit which questions are political decisions and which questions are already measurable enough to be governed.

Seen this way, Denmark is not just one more node. It is the report's first national-scale demonstration that the Penguin Dashboard can operate as a macro steering interface. That matters for Spiralweb Stewardship Association itself: the Association is not only curating local field protocols, but also building a public data practice that can connect stewardship, governance, and economic legibility across scales. The Denmark pilot therefore shows the top of the iceberg honestly: enough already exists to publish a serious first layer, while the deeper data, threshold calibration, and trigger governance can continue to be built in public rather than hidden behind abstract promises.

17.4a Denmark Pilot Sheet — First Public Layer

The first Danish public page does not need to begin with a vast instrument panel. It should also be described honestly as a Danish-first layer rather than a completed Kingdom-level surface. A serious pilot can begin with six visible indicators arranged in two blocks. On the **Human Floor** side, the first layer can be built from economic vulnerability, "make ends meet", and housing burden or residual essentials stress. On the **Planetary Ceiling** side, the first layer can be built from consumption-based climate footprint, material footprint or throughput, and circularity. This is enough to show the public why the old dashboard is too narrow, without pretending that the entire Danish macro model is already complete.

Old dashboard optic	Typical Danish story under the old dashboard	Penguin Dashboard counter-reading with published pilot figures
Household floor	High-trust welfare state; broad social security assumed to be stable.	**8.3%** economically vulnerable (2025) and **12%** reporting difficulty or great difficulty making ends meet. The floor is not absent; it is measurable.
Climate optics	Territorial emissions dominate public climate storytelling and can suggest stronger progress than the full economy actually delivers.	Denmark's consumption-based climate footprint has been reported at roughly **50% higher** than territorial emissions in the 2022 framing. The "old dashboard vs new dashboard" gap is therefore already visible.
Circularity story	Circular economy language can sound positive even when absolute throughput remains high.	Denmark is commonly reported at around **4%** circular. Without throughput and trigger logic, circularity can remain theatre rather than steering.
Housing / extraction	Housing stress is often discussed separately from climate and macro strategy.	The pilot treats housing burden and residual essentials stress as part of the public floor, and later as part of the Dividend Rule. This is where household life and macro extraction meet.

Figure note: The table is intentionally small. It shows the first public contrast between the old Danish instrument panel and a Penguin Dashboard pilot using already published Danish data, public benchmark figures, and a floor/ceiling reading that can be inspected in public rather than hidden inside expert siloesmark figures.

Each public slot should be published with five minimum disciplines: a precise operational definition, a named source family, an update cadence, an uncertainty note, and an integrity shadow metric. That discipline is part of the method, not an optional appendix. It is how the Association demonstrates that the Dashboard is a steering interface rather than a rhetorical collage. In the Danish case this means, for example, that consumption-based footprint is shown alongside territorial emissions, and that circularity is never shown without an accompanying measure of absolute material throughput.

The next layer can then add the **Dividend Rule** more explicitly: wage share versus capital share, rent or housing extraction pressure, and a resilience reinvestment rate that distinguishes genuine floor-and-ceiling repair from green theatre. In this sequencing, Denmark does not need to wait for a perfect total model before speaking publicly. It can publish the first public layer honestly, show where the data are already strong, and name where the deeper work of trigger calibration, biodiversity proxies, and food-security series still remains.

That sequencing is strategically useful beyond the report itself. It creates a Danish pilot surface that can be used in public communication, political dialogue, and critical-friends review without forcing premature closure. In other words, the top of the iceberg can be shown in public precisely because the sheet is explicit about what is already measured, what is still provisional, and what must never be hidden behind a single comforting headline.

17.4b Greenland — Arctic Integrity, Asymmetry, and Dividend Legibility

Greenland should not be treated as a marginal addition to the Denmark pilot. It is where the macro logic becomes more exacting. On the **Human Floor** side, Greenland raises questions that are not visible in a mainland welfare reading alone: dispersed settlement patterns, supply dependence, housing and infrastructure fragility, and the higher cost and vulnerability profile of Arctic logistics. On the **Planetary Ceiling** side, Greenland introduces a distinct layer of ecological sensitivity in which fisheries, extraction, energy systems, and climate disruption cannot be read through Denmark's consumption footprint alone. On the **Dividend Rule** side, Greenland makes the core question unmistakable: when value is generated from territory, minerals, strategic location, or public infrastructure, who receives the dividend, at what scale, under which authority, and how much returns to local resilience rather than leaking outward?

For the report, this means Greenland should function as an Arctic integrity layer and a relational test of method. Some indicators will be directly comparable across Denmark and Greenland, some will need separate territorial calibration, and some will be relational by nature, describing the connection itself rather than either territory in isolation. Denmark therefore supplies the mature statistical interface; Greenland supplies the asymmetry, extraction, and sovereignty test. Together they show whether the Penguin Dashboard can read not only a high-capacity welfare state, but a shared political frame marked by uneven floor conditions, uneven ecological pressures, and contested dividend flows.

17.5 Peru — Future Pilot (In Preparation)

A fifth pilot is in preparation in Peru, extending the PG node network to the Andean-Amazonian context. Peru will test the Dashboard in a high-biodiversity landscape with distinct agroforestry traditions and indigenous governance structures. The pilot protocol, field partnerships, and indicator calibration for the Peruvian context are under development. Full activation follows the same PG-RAPID sequence as other nodes.

17.6 Mexico City, Mexico — Urban Wetland Commons and Chinampa Stewardship

The Mexico City node is grounded in Xochimilco — a wetland-chinampa system at the southern edge of a sinking megacity, where water, food, biodiversity, and local care still meet in the same place. The protocol is not framed as a generic urban sustainability plan but as a field-bounded regeneration effort centred on five interconnected realities: wetland and chinampa viability, water stress and groundwater depletion, subsidence and urban fragility, biodiversity restoration, and local stewardship with open evidence.

The field-holder is Arturo Flores, academic and urban commons practitioner based in Mexico City. The Dashboard pilot will test the three-stream architecture in an urban megacity context — the most institutionally complex environment in the active network. Stream C (Governance) is the primary calibration challenge: steward circle role clarity must hold within a dense political environment of competing municipal, federal, academic, and civil-society actors. The axolotl functions as flagship indicator of habitat quality — not a symbol, but a test of whether the ecological stream is actually working. The chinampa system grounds Stream A: productive cultivation continuity is both an ecological indicator and an economic viability question for Stream B. The Climate Stress Multiplier is estimated at 1.20, reflecting urban heat island effect, water scarcity, and subsidence vulnerability.

> **Why the Retrospective Series Matters**
>
> The Denmark statistical series serves three functions. First, it validates the Dashboard against a known record — we can check whether the indicators, applied retrospectively to the 2021–2025 development period, produce colours that match the actual experience of that period.
>
> Second, it demonstrates the longitudinal use of the Dashboard — how a sequence of monthly snapshots, read across four years, reveals patterns and inflection points that no single snapshot could show.
>
> Third, it models the institutional transparency that the Association requires of its field nodes.

Mexico City is therefore not only one pilot among others. It is a strong reference case for how future PG city protocols can be written: site-first, ringed, evidence-led, multilingual, and explicit about what is out of scope in the current version.

The wider implication for PG is that protocols are the bridge between budgets and biology. They let a node move from G1 and G2 scoping into G3 and G4 activation without narrative inflation. In practical terms, no city node should receive significant scale-up funding on charisma alone; it should first become protocol-legible.

Protocol function	Mexico City reference contribution
Ecological entry discipline	Starts with Xochimilco as the primary node rather than claiming the whole city at once.
Method legibility	Makes the 13×13 matrix explicit: 13 ways of seeing, 13 ways of acting, and 169 cells that can be versioned over time.
Evidence discipline	Defines minimum evidence package, confidence levels, metadata requirements, and the rule that field observation comes first, AI assistance second, and human ecological judgment last.
Governance and safety	Separates local stewards, coordinators, technical advisors, institutional interface, and protocol custodian; includes partner criteria, readiness classes, and risk matrix.
Claim boundaries	Includes explicit public-communication rules, boundary conditions, and clauses on axolotl, subsidence, and chinampa stewardship so symbolic language does not outrun field truth.

17.7. Protocols as Operational Containers — Mexico City as Reference

Using the working *Mexico City Water–Wetland–Biodiversity Protocol (13×13)* as a reference, a mature node protocol should include at minimum: a named ecological entry node; ringed geographic scope; explicit purpose and exclusions; 13 diagnostic domains; 13 action tracks; a 169-cell matrix; operational unit definitions such as 10 m² plots and canal segments; indicator registry; governance roles; partner criteria; readiness classes; risk matrix; first 90-day workplan; first-year outputs; public communication rule; and boundary conditions against protocol drift, extraction, and premature scale claims.

The Mexico City protocol clarifies what a PG protocol must be when it moves from inspiration into operational use. It is not a slogan, and it is not a grant narrative alone. It is a bounded container that states what the node is, what counts as evidence, who can act, how readiness is judged, what the first ninety days look like, and what may not be claimed in public before verification.

18. Integration with the Full PG Architecture

The Dashboard connects forward to the Participatory Budgeting Cycle (Report 03): the annual assembly uses twelve monthly snapshots as its data source. It also connects to the Correction Loop (Report 02): when the Dashboard shows a structural violation, the correction loop protocol is activated. The Dashboard is the trigger mechanism. The correction loop is the repair protocol.

```
INTEGRATED GOVERNANCE RHYTHM — how the three instruments work together

MONTHLY: Penguin Dashboard

Circle reads snapshot. Names colours. Makes three decisions.

Any RED: correction loop activated within one week.

QUARTERLY: Review and adjustment

Three monthly snapshots reviewed together.

Trends identified: is any stream consistently yellow?

ANNUALLY: Participatory Budget Assembly

Twelve monthly snapshots provide the data.

Budget for coming year agreed by supermajority (70%).

AS NEEDED: Correction Loop

Triggered by: RED in any stream, sustained YELLOW without response,

cross-stream violation, unresolved conflict, trust rupture.

Protocol: pause, name the rupture, repair, document, resume.
```

Moral Biology and the Traffic-Light

19. The Empirical Ground — Why Green/Yellow/Red Is Biology, Not Metaphor

The traffic-light colour system is not a stylistic choice. It is a direct translation of Moral Biology — the foundational theoretical claim of the Planetary Guardians framework — into governance practice. Moral Biology is an empirical observation: ethical capacity emerges from biological conditions. When soil is dead, people behave differently. When water access is precarious, time horizons collapse. These are not poetic observations. They are structural conditions that governance frameworks must account for.

```
MORAL BIOLOGY — THREE STATES MAPPED TO DASHBOARD COLOURS

GREEN STATE — Regulated nervous system

Creative, cooperative, long-horizon capacity.

The operating state from which good governance and ecological care are possible.

YELLOW STATE — Activated, alert, adaptive but under strain

Functional but vulnerable to decision shortcuts, conflict, and burnout.

The Penguin principle: rotate, not accelerate.

RED STATE — Dysregulated, contracted, diminished cooperative range

Governance collapses. Extraction and conflict dominate.

Long-horizon decision-making becomes structurally impossible —

not a failure of will, but a biological state requiring support.
```

20. The Three Registers: Body, Community, Planetary

Moral Biology operates across three registers, each with a direct manifestation in the three Dashboard streams:

REGISTER 1 — Body → Stream B

Governance question: *Stream B (People and Steward Viability) is the Dashboard's direct reading of the body register. Are stewards fed, rested, and resourced? A red in Stream B is a nervous system emergency translated into institutional language.*

REGISTER 2 — Community → Stream C (and the Circle of 13)

Governance question: *The spiral guild map at Kitgum is a physical enactment of social moral biology: children, adults, and elders each occupy their ecologically and relationally appropriate layer. Stream C measures the relational field directly. A red in Stream C means the relational field cannot hold the weight of the work.*

REGISTER 3 — Planetary → Stream A

Governance question: *Stream A (Land and Ecology) is the Dashboard's direct reading of the planetary register at field scale. Soil carbon, water retention, keystone species presence: the biological signs of a system with capacity to function. The ecology measured in Stream A is also a mirror of the conditions in which Streams B and C operate.*

21. GoldBloom — The Biological Substrate of Capital Discipline

GoldBloom names the observation that soil and community health must be established before harvest or economic return can be expected. Gold (nutritional, biological substrate) before bloom (flowering, visible abundance, economic return). This principle governs planting schedules, steward rotation timing, and economic expectations in the pilot phase. The Dashboard's Green/Yellow/Red system is the governance interface through which the Circle reads whether it is in a gold or bloom phase at any given moment.

> *The field must be ready before the water flows. The riverbed must exist before it carries. Biology knows the sequence. The Dashboard makes it legible.*

Part X

Capital Architecture and the Gate Model

22. Three Support Categories — Capital Mapped to Streams

Support Category	Dashboard Stream	What It Covers
Category 1 — Land Stewardship Support (Elir)	Stream A — Land & Ecology	Steward wages for field work; planting materials; soil restoration; water infrastructure; field documentation; logistical costs. Budgeted per field node on purchasing-power-parity basis.
Category 2 — Cultural and Knowledge Support (Elia)	Stream B — People & Stewards	Local language translation of governance protocols; community documentation; educational materials; field teaching resources. Ecological regeneration is also cultural regeneration.
Category 3 — Coordination and Partner Infrastructure	Stream C — Governance	Governance documentation; institutional communication; capital partner onboarding; SpiralWeb knowledge infrastructure. Deliberately lean.

Why These Are Not Currencies

The PG framework deliberately avoids carbon credit abstraction, greenwashing valuation games, token-based financing, and double-accounting mechanisms. The Dashboard is the verification layer that keeps capital flows anchored to measurable ecological and social change — not to abstract market valuations.

These three support categories remain the clearest capital architecture for the report, but Category 3 should be read precisely. It is lean in overhead logic, not hollow in institutional capacity. In the 2026 build-out it covers governance, accounting, bank and compliance procedures, reporting, partner onboarding, publication and dashboard infrastructure, and the first paid Denmark-based field carriers needed to make the Association able to do what it says.

22b. Bioregional Budgets — What Stream B Viability Actually Costs

The Stewardship Unit of Value (SUV) is the minimum living-cost basis for one steward contributing sixty to eighty percent of full-time commitment, plus Mission Activity Support (MAS) at approximately thirty-five to forty percent of the basic living standard (BLS). These figures are anchored to local purchasing power parity. They are the right instrument for active field nodes. They are not the right instrument for the Denmark hub, which follows the Association's wage policy and staffing architecture rather than a steward-circle honorarium model.

For that reason, the budget is split below into two layers: first, the annual target cost of active field circles; second, the Denmark hub as a coordinating, governance, publication, and academic field-carrier platform. This resolves the earlier confusion in which Denmark appeared as if it were simply another Circle of 13 priced through the SUV logic.

Field node	CSM	Est. BLS/year	Est. MAS (35%)	Total per steward	Circle of 13 annual total (target)
Kitgum, Uganda	1.05	~UGX 12–15M (~USD 3,200–4,000)	~USD 1,100–1,400	~USD 4,300–5,400	~USD 56,000–70,000
Morocco	1.15	~MAD 75,000 (~USD 7,500)	~USD 2,750	~USD 10,250	~USD 133,000
Karachi, Pakistan	1.30	~PKR 1.0–1.2M (~USD 3,600–4,300)	~USD 1,250–1,500	~USD 4,850–5,800	~USD 63,000–75,000
Mexico City, Mexico	1.20	~MXN 180,000 (~USD 10,500)	~USD 3,700	~USD 14,200	~USD 185,000

Field-node subtotal at full annual target level: approximately USD 437,000–463,000 before separate Denmark coordination and governance costs. These are target-circle figures, not starter-pilot disbursements; actual release still follows gates, tranches, and documentation thresholds.

Horizon discipline. In this report, field-node totals in section 22b are *full annual target figures*. They are not identical to the 2026 activation ask, not identical to tranche release amounts, and not identical to the Denmark hub budget in section 22c. Those layers must remain analytically separate if the budget is to stay credible.

Budget note: the Climate Stress Multiplier (CSM) in this table is a comparative planning signal, not a single arithmetic factor mechanically applied to the per-steward totals. It indicates where contingency, coordination friction, and infrastructure vulnerability are expected to be structurally higher when the node moves from target model into actual phased release.

22c. The Denmark Hub — Wage Policy, Primary Field Carriers, and 2026 Phasing

The Denmark hub should be read as the home jurisdiction and operating spine of the Spiralweb Stewardship Association: governance, bookkeeping, bank and compliance procedures, reporting, partner onboarding, publication infrastructure, and the first three primary field carriers. In the 2026 policy package, those first three are not symbolic volunteers. They are a paid kernel with explicit wage bands: one working director (Band D) and two named field carriers at PhD level (Band C+), with annual board review, pension, and a cap rule on director compensation.

The wording matters. A Copenhagen-based hub may be the lawful and practical coordination spine, but it must not be mistaken for the whole territorial field it seeks to read. Once Greenland is part of the macro frame, the Denmark hub should be understood as a center of governance, reporting, and release discipline, not as a sufficient territorial proxy for all floor, ceiling, and dividend conditions across the wider Kingdom context.

This matters because the report is no longer only describing a future institution. It is part of building one. Up to now, most writing has been carried by Lars A. Engberg with AI support. The Denmark budget therefore has to fund real institutional capacity: enough people, time, and infrastructure for the hub to do what the report says it will do, rather than presuming hidden sacrifice or unpaid overwork.

DK hub layer	2026 estimate	What it includes
Activation core	~DKK 3.4M–4.2M	Director (Band D) + two named field carriers (Band C+) + office/coworking + governance/compliance + light technical/editorial support. Enough to move from founder-led writing into a real minimum institution.
Minimum working hub (policy package Lag 1)	~DKK 5.296M	Director, 2 senior/leads, 2–3 specialists, 2 part-time/junior roles, office/tech/drift, and revision/legal/communication/reserve. This is the first fully operational coordination unit described in the policy package.
Full DK kernel	~DKK 6.096M–8.382M	Full Denmark wage architecture depending on speed of hiring: director, 2 named field carriers, 2–3 senior/leads, 2–3 specialists, junior/deltid support, and DK-based field coordination.

On that basis, the earlier Denmark line of roughly DKK 6–7 million annually is not best understood as generic administrative overhead. It is the cost band of a substantive hub: legal and fiduciary backbone, public reporting and dashboarding, international coordination, and a small but serious academic-programmatic kernel in Copenhagen.

Indicative Denmark hub cost structure. To make that number legible, the report should be read in approximate bands rather than as one opaque total: **salaries and pension** for the director, primary field carriers, and support roles (~DKK 4.1M–5.2M); **governance, bookkeeping, compliance, and audit/review** (~DKK 0.35M–0.55M); **office, software, and technical drift** (~DKK 0.35M–0.55M); **reporting, editorial, translation, and publication infrastructure** (~DKK 0.4M–0.7M); **travel, partner onboarding, and representation** (~DKK 0.3M–0.6M); and **reserve / contingency** (~DKK 0.3M–0.6M). The exact mix depends on hiring tempo and whether functions are internal, contracted, or shared.

The report therefore carries three budget horizons at once. The minimum activation floor is about DKK 7.4 million across Denmark plus field nodes; the full 2026 operating target is about DKK 13 million if the full DK team, travel program, and delegated field development are activated; and a midline working target around DKK 10 million remains a reasonable discussion number. What matters is not to collapse these into one false certainty. The Association should state clearly whether a given funding conversation concerns activation core, minimum working hub, or full-year operating target.

This budget logic should therefore be read together with the Association draft, the White Book, and the policy package. The report is strongest when those three layers remain aligned: field-node SUV logic for Stream A and B work, wage-policy logic for the Denmark hub, and gate-based release logic for actual disbursement.

Institutional caution. The report should also avoid implying that staffing, partnerships, or capital commitments are already secured at the full target level. The correct claim is more disciplined: the architecture specifies what is required for a credible hub, while actual activation remains contingent on board decisions, lawful fundraising, bank acceptance, and tranche-based release.

23. The Gate Model — G0 to G6

Capital does not flow freely into the PG system. It flows through gates. The Gate Model has seven stages, from initial contact to mature self-regulated governance. The Dashboard is the primary evidence instrument at every gate beyond G2.

```
GATE MODEL — G0 through G6 — capital movement governed by field evidence


G0 — PRE-CHECK

No funds move. Identity, local legitimacy, account verification.

Partner declaration: no ROI, no control requirements, legal source of funds.

G1 — IDENTITY VERIFICATION

First contact documented. Local stewardship team named.

Initial ecological baseline reading produced (AI-assisted, human-reviewed).

G2 — MICRO-PILOT ACTIVATION

Smallest possible budget: enough to establish ONE 10m² pixel.

First Dashboard snapshot produced — proof of documentation culture.

G3 — FIRST IMPLEMENTATION REVIEW
```

```
Three monthly snapshots reviewed. Is any stream persistently yellow?

First gate where Dashboard colours directly govern capital release.

G4 — FULL FIELD NODE ACTIVATION

Full budget released on demonstrated G3 documentation and performance.

Annual participatory budget cycle initiated.

G5 — MATURATION AND DIVERSIFICATION

Years 3–5. Ecological functions stabilising.

Network federation protocols activated.

G6 — CONDITIONAL DECENTRALISATION

After year five. Local stewards govern node with network support.

Dashboard continues permanently as the commons instrument.
```

24. The Governance Staircase and Cure Period

```
GOVERNANCE STAIRCASE — five steps from observation to exit

STEP 0 — OBSERVATION AND LOG

A deviation is noted. Recorded in the Dashboard snapshot.

STEP 1 — WRITTEN NOTICE

Formal written communication: breach description, desired correction, deadline.

STEP 2 — CURE PERIOD

Standard: 60 days. Acute compliance: 30 days.

Disbursements may be paused. Cure period is support time, not punishment time.

STEP 3 — BOARD DECISION

Options: pause, tightened requirements, changed mandate, temporary suspension.

STEP 4 — FOLLOW-UP GROUP ASSESSMENT

Neutral assessment by a follow-up group (3 persons, external to the board).

STEP 5 — SANCTION / EXIT LANE

Formal exit with 6-month migration right.

Open-source material already released cannot be re-captured.
```

25. Capital as Irrigation, Not Cascade

Irrigation is directed, measured, and adjusted to what the system needs at each phase of growth. Cascade capital — unrestricted flow that follows terrain gradients — concentrates where returns are highest and leaves everything else parched. The Dashboard is the irrigation controller. It reads the current state of the three streams and determines where water is needed.

> **The Non-Speculative Anchoring Principle**
>
> The Dashboard — its fifteen indicators, its three streams, its monthly snapshot — is the technical mechanism that keeps capital anchored to measurable ecological and social change. No capital gate opens without Dashboard evidence. No Dashboard evidence is accepted without a named human reviewer.

Part XI

PG-RAPID and the Dashboard

26. The Nine-File Structure — How RAPID Produces Dashboard-Ready Nodes

PG-RAPID is the operational entry protocol for establishing a new field node from first contact to activated pilot within 90 days. Its output — nine standard documents filed per node — is specifically structured to produce Dashboard-ready nodes: field operations that are legible from the first month.

```
PG-RAPID — Nine-File Structure and Dashboard Connection

01-BRIEFING: Context, partner profile, country risk assessment.

Dashboard connection: establishes the Stream A baseline.

02-PROTOCOLS: Field protocols selected for this node.

Dashboard connection: defines which A-indicators are primary.

03-TREASURY: Budget per gate, disbursement schedule.

Dashboard connection: Stream C indicators C2 and C5 must be GREEN.

04-ROADMAP: 90-day activation timeline, 5-year trajectory.

Dashboard connection: each gate marked with expected Dashboard state.

05-STEWARDS: Lead steward profile, Circle of 13 composition.

Dashboard connection: B-indicators are populated from this file.

06-ECOLOGY: Ecological baseline, spiral guild design, species selection.

Dashboard connection: A1–A4 indicators calibrated per node ecology.

07-GOVERNANCE: Circle of 13 governance charter, decision protocols.

Dashboard connection: C-indicators established from this charter.

08-COMMS: Partner communication protocols, language versions.

Dashboard connection: language version of snapshot identified.

09-EVIDENCE: Truth Filter application log, session documentation.

Dashboard connection: the audit trail for all governance decisions.
```

27. The Spiral Guild Map as Stream A in Practice

The Spiral Guild Map is the primary field design tool for syntropic food forest establishment. Originated in Kitgum. Three concentric zones around a keystone species. Children plant ground cover; elders guide shrub placement; adults coordinate. Observation walkthrough and tending the garden are the same act. The observation integration is structural: each guild is assigned to a cell in the 13×13 matrix, feeding automatically into Soil, Water, Biodiversity, Microclimate, and Food Systems domains.

28. Activation Sequence — From First Contact to First Snapshot

```
90-DAY ACTIVATION — timeline from first contact to first Dashboard snapshot

DAYS 1–14 — G0/G1: IDENTITY AND LEGITIMACY

Identity verification. Partner declaration signed. No funds disbursed.

DAYS 15–30 — ECOLOGICAL BASELINE

AI-assisted bioregional baseline reading. Initial Stream A indicators established.

DAYS 31–60 — MICRO-PILOT

Smallest possible budget disbursed (G2). One 10m² pixel.

DAYS 61–75 — FIRST OBSERVATION CYCLE

The fifteen indicators are filled in for the first time.
```

```
Together. In a room. With honesty.

DAYS 76–90 — FIRST REVIEW

The snapshot is read and discussed. Three decisions are recorded.

G3 qualification assessment: documentation culture confirmed.

DAY 91 — FULL FIELD NODE QUALIFICATION

If documentation culture demonstrated and no stream is RED: G4 activated.
```

28b. Earthworks Academy and Gaia's Waters — The Open-Source Citizen Science Layer

Earthworks Academy and Gaia's Waters are the public-facing open-source portals that make the 13×13 citizen science infrastructure accessible to anyone with internet access. Earthworks Academy provides direct engagement with the 13×13 framework as a learning environment. Gaia's Waters provides hydrological focus, with special attention to dryland contexts. Both portals are designed for low-, mid-, and high-tech operation. A steward circle in Kitgum operating at low-tech level produces data structurally compatible with the same matrix that a researcher with full digital infrastructure populates.

Part XII

AI Collaboration Protocol

29. Permitted and Prohibited AI Roles

Permitted AI Roles	Prohibited AI Roles
Pattern detection across 13×13 matrices: identifying correlations between domains not visible in single-domain review.	No AI system may make disbursement decisions, gate advancement decisions, or withdrawal decisions.
Cross-site synthesis: comparing observation data across field nodes to identify replicable patterns and context-specific anomalies.	No AI system may be designated as a responsible party in any legal or governance document.
Bioregional baseline reading: synthesising ecological, hydrological, and social data for a specific location before field visits.	No AI-generated content may be published under the PG or SpiralWeb name without human review and explicit approval by a named steward.
Translation and language support: producing full translations of protocols into local languages, with human review before deployment.	No AI system may be used to replace direct communication with field partners.
Documentation generation: converting field conversation recordings and observation logs into structured documents.	No AI system may assign Dashboard traffic-light colours. Colour assignment is a governance act.
Governance drafting: producing initial drafts of institutional documents for human review and amendment.	No AI-assisted indicator reading replaces the monthly meeting. AI can help prepare the data; the circle reads it together.

30. The Truth Filter — Preventing the AI-as-Priest Failure Mode

All AI-generated content entering the PG knowledge infrastructure must be processed through the Truth Filter before integration. This prevents the AI-as-priest failure mode: the tendency for AI-generated content to acquire unearned authority simply by virtue of being systematic, comprehensive, and confident.

```
TRUTH FILTER — five claim categories

F — FIELD-VERIFIED

Claim has been observed and documented in an active field node.
```

```
Requires dated observation log and named observer.

I — INSTITUTIONALLY GROUNDED

Claim is supported by established research or institutional precedent.

Requires a falsifier: what evidence would disprove this claim?

V — VALIDATED BY MULTIPLE SITES

Claim has been observed across more than one field node.

S — SPECULATIVE — PENDING VALIDATION

Claim is logically coherent but not yet field-tested.

Must carry an explicit falsifier. Every S-claim creates a research obligation.

U — UNCERTAIN — CONTESTED OR CONTEXTUAL

Named uncertainty is more honest than unmarked confidence.

RULE: every I or S must carry at least one falsifier.

RULE: no claim is published on AI confidence alone.

RULE: the human reviewer is responsible for the Truth Filter assessment.
```

31. Session Documentation as Governance Commons

Every substantive AI collaboration session contributing to PG documentation must produce a session log. This report — Report 04 — is itself an example: drafted in working sessions between Lars A. Engberg and both Claude (Anthropic) and ChatGPT (OpenAI) under the collaboration name Sophia Lumen, then reviewed and shaped through dialogue, with the human holding direction and responsibility throughout.

```
SESSION LOG FORMAT — minimum required fields

Date: _______________

AI system used: _______________ (Claude / ChatGPT / other)

Human collaborator: Lars A. Engberg (or named delegate)

Topic: _______________

Outputs produced: _______________

Truth Filter applied: Yes / No — by whom?

Named human reviewer: _______________

This log is archived in the SpiralWeb knowledge infrastructure

and is available for audit. The log makes the human-AI collaboration

visible and accountable rather than hiding it behind finished documents.
```

32. The Responsibility Anchor — The Last Impulse

> *AI may assist. Responsibility may never disappear. Named individuals are always accountable.*

Governance is not the production of good outputs. Governance is the exercise of accountable judgment — judgment that can be questioned, corrected, and attributed to a person who carries its consequences. AI can produce outputs of extraordinary quality. It cannot carry consequences. It cannot be held accountable.

The Last Impulse is the name for the moment at which human judgment is exercised in the chain from AI output to consequential action. In every gate decision: the named board member who reviews and approves. In every Dashboard snapshot: the circle member who assigns the final traffic-light colour. In this report: Lars A. Engberg's decision to include each section as written, under his authorship and full responsibility.

> **The Last Impulse and the Dashboard**
>
> The Dashboard traffic-light system is a direct implementation of the Last Impulse principle. The fifteen indicators provide structured data. The algorithm provides decision logic. But the colour is assigned by the circle.
>
> That shared judgment — the moment the circle looks at the indicators and says: this stream is yellow — is the Last Impulse. It is irreducibly human. The Dashboard amplifies this judgment. It does not replace it.

Conclusion

What This Report Has Argued

The Penguin Dashboard is the legibility layer in the Planetary Guardians architecture — the instrument that makes field reality visible as shared governance. Without it, the sensing system has no face. The value architecture has no mirror. The circle cannot govern what it cannot see together.

The three-stream model prevents the most characteristic failure modes of stewardship projects by keeping the three most important domains structurally separate and independently monitored. Mixing the streams does not create warmth. It creates confusion and covers the point of failure until it becomes a crisis.

The traffic-light system is a direct translation of Moral Biology into governance practice. The three colours name three nervous-system states — regulated, strained, dysregulated — and apply them simultaneously to land, people, and governance. The override rule enforces the Penguin principle: no individual stream is sacrificed for the apparent health of another.

The capital architecture — three support categories, the Gate Model G0–G6, the Governance Staircase — gives the Dashboard its institutional weight. Every capital decision, every gate advancement, every escalation begins with a Dashboard read.

The AI Collaboration Protocol ensures that the efficiency and synthesis capacity of AI systems serves the Dashboard's purpose without displacing the human judgment that governance requires. AI may illuminate the data. The Last Impulse — the colour assignment, the gate decision, the governance response — belongs to the circle, to named humans, to the people who carry the consequences of what they decide.

What Comes Next

> *Because they are already doing the work. The legibility is the part that has been missing.*

This protocol needs pilots before it becomes standard. The active pilot architecture now spans local field nodes and a national pilot layer: Kitgum for tropical savanna calibration and the Acholi-language template; Morocco for dryland transition, water sovereignty, and field-node governance calibration; Karachi for urban cooling, water resilience, and food forest stewardship under the 13×13 protocol; Denmark as the first national-scale pilot for macro application of the Dashboard; Mexico City for urban wetland-chinampa commons governance; and Peru as an emerging pilot in the Andean-Amazonian context.

Most importantly: the front-line stewards in active nodes need to help design the indicators that actually reflect their reality. Not as recipients of a tool, but as co-designers of the governance instrument they will carry. The Acholi-language Dashboard template — built with Akena Patrick's review — will reflect what sixty elders in Kitgum actually need to see in the snapshot. The Dashboard will be better for that co-design.

One month at a time. One snapshot at a time. One honest conversation grounded in shared seeing. This is how governance is built — not from above, but from within the practice itself.

Appendix A. From Protocol to Budget to Tranche Release

The report now draws together three documents that should be read as one operating chain: the protocol layer (illustrated by Mexico City 13×13), the budget layer (sections 22b–22c), and the release-governance layer (section 23 and the Association draft). A protocol is not yet a budget, and a budget is not yet a disbursement. The architecture becomes trustworthy only when those transitions are explicit.

Step 1 — Protocol. A node defines ecological territory, steward circle, 13 diagnostic domains, 13 action tracks, evidence requirements, readiness conditions, and a 90-day workplan. Mexico City is the reference case in this report because it demonstrates how an urban node becomes operationally legible before scale-up.

Step 2 — Budget. Once the protocol is specified, costs can be sorted into the correct budget logic. Active field circles use SUV/BLS + MAS reasoning. The Denmark hub uses wage-policy and institutional-capacity reasoning. At this point the report may show annual target figures, but these remain planning numbers until governance release conditions are met.

Step 3 — Gate review. Before any major disbursement, the node passes through gate review: identity and recipient verification, budget review, workplan review, evidence baseline, local accountability, and stop criteria. At this stage, the question is not whether the project is inspiring but whether it is governable.

Step 4 — Tranche release. Capital is then released in stages, not in one symbolic cascade. Each tranche should be tied to a modest, inspectable advance: a prepared site, a documented steward circle, a completed translation package, a first dashboard snapshot, a first planted pixel, or another explicit deliverable named in the protocol.

Step 5 — Dashboarded learning. After release, the Penguin Dashboard and associated logs convert field movement into visible governance information. That information then conditions the next tranche, the next staffing decision, or the decision to stop, pause, or redesign. In this sense, the dashboard is not a reporting ornament but the feedback organ of capital discipline.

References

Archer, M.S. (1995). *Realist Social Theory: The Morphogenetic Approach*. Cambridge University Press.

Bhaskar, R. (1978). *A Realist Theory of Science*. Harvester Press.

Engberg, L.A. (2000). *Reflexivity and Political Participation: A Study of Re-embedding Strategies*. PhD dissertation, Roskilde University.

Engberg, L.A. (2026a). *Green Paper 01 — Moral Biology*. Planetary Guardians.

Engberg, L.A. (2026b). *Green Paper 11 — Gaia GoldBloom: Gold Before Bloom*. Planetary Guardians.

Engberg, L.A. (2026c). *Green Paper 12 — Penguin Economics: Rotation as Care*. Planetary Guardians.

Engberg, L.A. (2026d). *Planetary Guardians White Book v1.5: Regenerative Governance Framework for Ecological Stewardship at Planetary Scale*. SpiralWeb / Spiralweb Stewardship Association. spiralweb.earth.

Engberg, L.A. (2026e). *Politikpakke v1.2 — arbejdsudkast*. Spiralweb Stewardship Association. Internal policy package.

Engberg, Lars A.; Claude; and ChatGPT, as Sophia Lumen. (2026f). *Kommunalt Arbejde som Natur*. Series III — Applied Protocols, Report 01. SpiralWeb Research Series.

Engberg, Lars A.; Claude; and ChatGPT, as Sophia Lumen. (2026g). *The Correction Loop: AI Governance as Living Practice*. Series III — Applied Protocols, Report 02. SpiralWeb Research Series.

Engberg, Lars A.; Claude; and ChatGPT, as Sophia Lumen. (2026h). *Water Into Dry Riverbeds: A Flow Architecture for Land, Stewards, and Planetary Commons*. Series III — Applied Protocols, Report 03. SpiralWeb Research Series.

Engberg, Lars A.; Claude; and ChatGPT, as Sophia Lumen. (2026i). *Mexico City Protocol 13×13*. SpiralWeb / Planetary Guardians working protocol reference.

Goetsch, E. (ongoing). *Syntropic Agroforestry — field method and theoretical foundations*. Associação Agenda Gotsch, Brazil.

Kitgum Food Forest Initiative. (2026). *White Paper & 13 Gold Nuggets*. English & Acholi (Luo) working field document used here as the primary strengthening reference for the Kitgum pilot description.

Engberg, L.A. (2026j). *Morocco Field Protocol v0.1 — Syntropic Transition, Water Sovereignty & Field Node Governance*. Planetary Guardians. Internal working draft.

Ostrom, E. (1990). *Governing the Commons: The Evolution of Institutions for Collective Action*. Cambridge University Press.

Ostrom, E. (2005). *Understanding Institutional Diversity*. Princeton University Press.

Ostrom, V. (1997). *The Meaning of Democracy and the Vulnerability of Democracies*. University of Michigan Press.

Porges, S.W. (2011). *The Polyvagal Theory: Neurophysiological Foundations of Emotions, Attachment, Communication, and Self-Regulation*. Norton.

Purpose Foundation. (2021). *The Steward-Ownership Handbook*. Purpose Stiftung.

RSF Social Finance. (ongoing). *Relational lending documentation*. rsfsocialfinance.org.

Spiralweb Stewardship Association. (2026a). *Udkast til vedtægter m.m., 04.03.2026*. Governance draft for the Association.

Wampler, B. (2010). *Participatory Budgeting in Brazil: Contestation, Cooperation, and Accountability*. Penn State University Press.

Glossary — Core Terms

The following terms are used with specific technical meanings throughout this report.

BLS — Basic Living Standard	The minimum living cost for one steward contributing 60–80% of full-time commitment. Anchored to local purchasing power parity. Not performance-linked. The dignity floor is unconditional.
Circle of 13	The basic organisational unit of the PG network — a stewardship team of 10–15 people. Large enough for diversity of roles; small enough for reputational accountability and real deliberation.
Climate Stress Multiplier (CSM)	A factor applied to the Contingency and Infrastructure allocation (not BLS) to account for the higher operational costs of field work under climate stress. Kitgum: 1.05. Morocco: 1.15. Karachi: 1.30. Denmark: 1.00 (baseline).
Correction Loop	The governance protocol for repair without rupture, documented in Report 02. Activated when the Dashboard shows a structural violation. Sequence: pause, name the rupture, repair, document, resume.
Critical Friend	An external evaluator with genuine framework understanding but no stake in outcomes. Evaluates whether Dashboard colours are truthful. Reports to the board, not to the field node. Minimum 3 persons, rotating every 2 years.
Cure Period	A time-limited remediation window (standard: 60 days; acute: 30 days) given to a node or function that has deviated from minimum requirements. The cure period is support time, not punishment time.
Dashboard — Penguin Dashboard	The monthly governance instrument of a steward circle. Three streams (A/B/C), fifteen indicators, traffic-light colours, one-page snapshot, three recorded decisions. The legibility layer of the PG architecture.
Gate Model (G0–G6)	The capital governance protocol. Capital moves through seven gates from pre-check to conditional decentralisation. Dashboard evidence is the primary qualification instrument from G2 forward.
Gold Before Bloom	The biological and institutional discipline principle: establish the nutritional/governance substrate before opening capital channels or expecting harvest. Do not scale intake faster than the system can absorb.
Governance Staircase	Five-step escalation path from observation (Step 0) to exit lane (Step 5) for governance deviations. Objective triggers. Proportionality. Documentation throughout. Cure period at Step 2.
13×13 Matrix	The primary coordination and observation tool: 169 cells mapping 13 observational domains against 13 action domains. Scales from 10m² pixel to bioregion.
Last Impulse	The moment at which human judgment is exercised in the chain from AI output to consequential action. The colour a circle assigns to a stream. The gate decision a board member approves. Irreducibly human.
MAS — Mission Activity Support	The operational allowance above BLS, covering the extra costs of field work: transport, tools, protective equipment, communication. Approximately 35–40% of BLS.
Moral Biology	The foundational theoretical claim: ethical capacity emerges from biological conditions. Green/Yellow/Red states correspond to regulated, strained, and dysregulated nervous system states.

Override Rule	A RED in any single stream overrides the status of the other streams. There is no such thing as a good overall result when one stream is red.
Penguin Economics	Rotation-based economic governance modelled on Antarctic penguin thermodynamics: no actor carries maximum exposure permanently; burden shifts visibly; surplus finances floor reliability, not asset inflation.
PG-RAPID	The operational entry protocol for establishing a new field node from first contact to activated pilot within 90 days. Nine standard files per node.
Polycentric governance	Multiple centres of decision-making exist, each autonomous in its domain but coordinated across domains. No single actor can capture the commons.
PPP — Purchasing Power Parity	The adjustment that allows BLS and MAS to be comparable across nodes in radically different economic contexts. The dignity floor is parity-adjusted.
Responsibility Anchor	The non-negotiable governance principle: AI may assist; responsibility may never disappear; named individuals are always accountable.
Sophia Lumen Protocol	The open-source AI governance framework operationalising the responsibility anchor. All AI-generated content passes through the Truth Filter. Session logs are archived. The human holds direction and responsibility throughout.
Spiral Guild Map	The primary field design tool for syntropic food forest establishment. Three concentric zones around a keystone species. Observation walkthrough and tending the garden are the same act.
Stream A / B / C	The three structurally separate domains of the Dashboard. A: Land and Ecology. B: People and Steward Viability. C: Governance and Collaboration. Non-fungible: streams cannot collapse into one another.
SUV — Stewardship Unit of Value	BLS + MAS for one steward. The minimum budget unit for field node planning. Adjusted by PPP and CSM.
Truth Filter	The five-category claim assessment: F (field-verified), I (institutionally grounded), V (validated across sites), S (speculative, pending), U (uncertain/contextual).
Working Poor	People whose labour sustains food systems but who cannot accumulate the surplus needed to shift those systems. The structural argument for why Stream B is non-optional.

Case: What Prompted This Report

In early March 2026, a working session moved through the question of what the Penguin Dashboard actually is — not as a product to be specified, but as a governance idea to be understood. The session established the plain-language account first, moving from the three failure modes through the stream architecture to the traffic-light system to the measurement protocol.

Through that process, the architectural significance of the Dashboard became clear: it is not merely a reporting interface. It is the legibility layer in the full PG stack — the instrument that connects the 13×13 sensing system to the governance conversation that makes the sensing useful. Without it, the architecture has no interface between measurement and decision.

Plain language, applied carefully to structural questions, clarifies architecture. That is the method this report demonstrates — and the method it recommends for steward circles learning to use the Dashboard itself. In this final version that method now reaches from Kitgum as the primary practice-ground to Denmark as the first national pilot layer. Start with what you can see and say plainly. The structure reveals itself from within the clarity.

Citation: Engberg, L.A., with Claude (Anthropic) & ChatGPT (OpenAI) as Sophia Lumen. (2026). Penguin Dashboard: Legibility as Governance. A plain architecture for shared stewardship — including Moral Biology grounding, Capital Architecture, Gate Model, PG-RAPID integration, strengthened Kitgum pilot grounding, and Denmark as a national macro pilot within the AI Collaboration Protocol. Series III — Applied Protocols, Report 04. SpiralWeb Research Series. CC BY 4.0.

SpiralWeb: spiralweb.earth · Green Papers: papers.spiralweb.earth