



Turning fragmented project evidence into auditable, funder-ready learning and reporting

Insights from designing a Notion-based monitoring layer for the Agroecological Regenerative Cacao (ARC) project (2021–2026)

CIAT

Diego Paredes, MEL Regional Expert
Yovita Ivanova, ARC Project Manager
Project website: <https://allbiociat.org/44JEZ3C>

EXTERNAL ISSUES

July 2026

Key messages



The real reporting challenge in multi-partner projects is rarely a lack of work; it is typically a lack of standardization. When evidence is fragmented, accountability becomes expensive.



A lightweight relational structure can change that equation. The results integration layer, built on Notion's no-code workspace, treats evidence not as a filing system but as a connective backbone, linking what projects do every day to what institutions must report.



The value of the results integration layer lies not in the platform itself, but in the discipline it enforces, reaching agreement up front on what success looks like, how it will be tracked, and what counts as acceptable evidence. The tool only works when those decisions have already been made.



Incentives matter. When the same system that satisfies institutional reporting also provides project managers real-time visibility into their work, maintaining it stops feeling like a compliance task and becomes a useful habit.



Partners will not all adopt the same tools, and many should not be expected to. What matters is not conformity during data collection, but translation upstream into a shared evidence structure.

Context: a familiar scene, a different ending

It often starts the same way. A reporting deadline approaches, a funder meeting is scheduled, and the same questions surface. What exactly was delivered in each country? Which activities contributed to which outcomes? Where is the evidence that an indicator moved, and who can verify it?

In a research-for-development network like the Alliance's, those questions become harder to answer—not because projects are inactive, but because evidence is housed across different domains. Sub-grantees will not have access to institutional repositories or use Microsoft SharePoint, and may rely on Google Drive and local systems by necessity. Over time, projects accumulate a parallel universe of documents, including a very well-sorted list of resources (activity narratives, attendance lists, technical memos, training materials, partner agreements, etc.). All perfectly meaningful locally, yet difficult to assemble into an institutional reporting logic without extensive manual work.

Over the final year of the Fonds Francais pour l'Environnement Mondial-funded Agroecological Regenerative Cacao (ARC) project, operating across Colombia, Ecuador, and Peru, a structured evidence layer was built in Notion by the Alliance Monitoring, Evaluation, and Learning (MEL) Specialist, Diego Paredes, working closely with the project manager, Yovita Ivanova. The result was not a new database for its own sake. It was a bridge, a way of absorbing diverse evidence formats while preserving the relationships that make results intelligible, traceable, and defensible.

How the system was built and how it works

The problem was familiar before the solution existed. Every year, the ARC project produced a detailed technical report presenting components, results, activities, indicators, baselines, and targets, and, for each advance, a knowledge product with a direct link to supporting evidence. That structure was already understood by all the partners. It was the institutional language everyone spoke. But what if that same structure became the backbone of the monitoring system, instead of a separate document assembled once a year?

That is how the results integration layer began. Not with software selection, but with two source documents. On one side, the project proposal, which defined what had to be achieved, and the annual technical reporting template, which defined how progress was advancing and had to be demonstrated. We read both documents, mapped the logic, and built a Notion workspace to mirror it exactly. Partners did not need to learn a new reporting framework.

The planning area of the integration layer maps each activity in the work plan, with its key performance indicator (KPI) targets, responsible partner(s), geography, and timeline. Meanwhile, the reporting comes from the accumulation of knowledge products as project implementation unfolds, such as field reports, training records, datasets, technical memos, certification documents, and partner outputs that demonstrate field work and key milestones. Every time one of these resources is mapped, it must be linked to the activities it substantiates, thereby contributing to its KPI. That link is what transforms a document library into an accountability trail.

Because Notion's relational logic automatically pulls linked records together (Figure 1), any activity can be queried to show not just whether it was completed, but also by whom, where, when, and with what evidence, across multiple years of implementation. A project leader can move easily from an indicator figure to the activities that drove it, and from there to the specific documents that prove it. The accountability is already in the record, and any project that reports against a structured, logical framework can replicate it.

The architecture is intentionally simple. Three linked datasets form the backbone, which together formed these results for the ARC project:

- ✓ **Activities:** the work plan and accountability structure.
- ✓ **Knowledge products:** 159 documented items linked to indicators, with evidence-linked URLs, authors, lead institutions, partner organizations, and product type all captured as metadata.
- ✓ **Capacity development:** 68 events logged, including participation, topic, geography, and stakeholder group details recorded by group, not individually, enabling aggregation across the full project life cycle.




Figure 1

Schema for information flows: from fragmented evidence sources through the Notion bridge layer to institutional and funder reporting outputs.

Information translation upstream is the operative principle of the Notion-based system. Information flows from the project leader and the technical specialist to the MEL specialist, who enters and links it within the system. Partners continue using their own tools (Google Drive, local files, cooperative records). The MEL layer absorbs, integrates, and connects the project outputs into a shared structure. The ARC project's network spans 40 organizations headquartered across 10 countries (Figure 2), each engaged at different levels of intensity. Some are active implementers driving day-to-day activities; others participate consequentially, with their involvement defined by a specific function (certification, institutional partnerships, or strategic alliances). Together, they form a heterogeneous but coherent network of actors.

Insight Note

Turning fragmented project evidence into auditable, funder-ready learning and reporting

Partner Organizations Network

47 organizations across 10 countries

**Figure 2** ARC project partner organizations network, by headquarters locations.

That rich stakeholder diversity is precisely the challenge this system was designed to address. No single tool could be imposed. Instead, the system required a clear minimum standard at the translation point: any evidence entering Notion needed to include the author, institution, date, product type, and a link to the source document. Everything else was flexible.

The results

The ARC project is approaching its completion after five years. Below is the inventory that the Notion system made legible and reportable. The jump from 56% to 95%

reporting completeness was not achieved by doing new work, but by making existing work findable, linkable, and auditable. Completeness here refers to the proportion of activities with KPI targets met, traceable to substantiating evidence, and measurable, as the metadata structure automatically counts against those KPIs.

The documentation exercise provided evidence of how the project connected 2,070+ producers across three countries to new certifications, tripled yields in Colombia, and drove record exports in Peru, all while placing 750 ha under conservation.

Project Results

Key outcomes across environmental and socio-economic dimensions

750 ha

Under conservation agreements
(3 countries)

2,070

Producers trained on
environmental topics

1,198

Producers with new Rainforest
Alliance certification

Environmental impacts

Land restoration, biodiversity & climate adaptation

- **516 ha** of new agroforestry cacao systems (SAF) established across Colombia, Ecuador, and Peru
- **266 ha** of existing hectares renovated with agroforestry systems in Ecuador and Colombia for fine flavor cacao
- **704.8 ha** additionally rehabilitated through complementary soil management practices
- **Soil fertility** improved in Peru: composite scores rose from 8.4 to 9.0 in established SAF plots
- Deforestation monitoring systems deployed
- Cacao traceability systems in place (**3 cooperatives**)
- **3 conservation plans** approved by cooperatives
- **750 ha** under conservation agreements across 3 countries (Peru: **160 ha**; Ecuador: **36 ha**; Colombia: **444 ha**)
- Biodiversity (flora and fauna) increased in SAF plots
- **43 new native cacao genotypes** identified and characterized (contributing to a 60-accession clonal garden)
- Climate change adaptation through clone selection, propagation methods, SAF, soil restoration, and deforestation control
- Intra-parcel deforestation in Peru declined from **39.4 ha** (2021) to **13.9 ha** (2024)

Socio-economic impacts

Certification, trade & producer income

- **<5,000 ha certified** (organic / Fairtrade / Fair for Life)
- **1 new Rainforest Alliance certification** covering 5,373 ha and 1,198 producers
- **1 new commercial alliance** in Colombia — first **60 metric tons** of organic cacao exported to France (KAOKA)
- Organic volumes in ASOCAPIC grew from **zero** (2022) to **14,186 kg** by 2025
- **FOB price premiums of 20-25%** per kg linked to organic certification
- Yields in Colombia nearly tripled from **134 to 396 kg/ha/year** (2023-2025)
- Average farmer **income rose 87%**, from **\$1,410 to \$2,639 USD**
- Ecuador (CECAO) reached a record **2,070.4 tons** in 2025 — highest annual volume on record
- Peru (Colpa de Loros) export volumes grew from **304,482 kg (\$934K)** in 2021 to **562,358 kg (\$5.6M)** in 2025
- **1 organic fertilizer plant** established
- Traceability systems (KoboCollect) operational across all 3 cooperatives, aligned with EUDR requirements

Figure 3

ARC project results summary across environmental and socio-economic dimensions for 2021-2026.

Insight Note

Turning fragmented project evidence into auditable, funder-ready learning and reporting

Behind those numbers lay a deliberate sequencing, soil fertility scores rising in Peru's agroforestry system plots, yields climbing from 134 to 396 kg/ha in Colombia, and intra-parcel deforestation falling by 65%. The kind of foundational work that gave cooperatives the credibility to negotiate with international buyers and the organizational capacity to hold that ground once project support receded.

That evidence base translated directly into the Alliance's institutional reporting system (STAR) as outcome stories (OICR), innovations, and capacity development activities

books, which are visible in the Alliance of Bioversity and CIAT [results dashboard](#) and the broader CGIAR [bilateral portfolio](#) reach.

All have been documented in the Alliance OICRs in separate stories. These reports, available in the [institutional repository](#), capture both early shifts in (i) farmer and cooperative behavior and deeper, (ii) documented changes in land use, (iii) market access and income across three countries.

Regenerative Agroforestry Rehabilitates Amazon Cacao



Conserving Native Cacao Genetic Diversity



Strengthening Zero-Deforestation Cacao Supply Chains



The system is used, not just there to be presented

One question worth addressing directly is whether the system was actually consulted between presentations, or whether it only comes alive in front of a funder? The system's usage analytics provide a clear answer. Over ten months, the Notion workspace accumulated 320 views, with a visible spike cluster at launch in June and July 2025, a secondary pulse in September, and persistent traffic through to the project's close.

The implications are significant. Funders and counterparts are regularly accessing the system independently, without guidance. It functions as a live reference, not as a curated presentation. In quarterly follow-up meetings with the

funder, the Notion workspace has been on the agenda and has been open and navigable in real time. That kind of institutional credibility emerges from a system that is consistently maintained and genuinely useful.

“

It helps keep discussions focused, supports credibility, and reduces the time cost of answering follow-up questions

”

Yovita Ivanova, Project Coordinator

What made it operational?

Being realistic about conditions. The system was not designed for an ideal world of rigorous partner compliance and unified data infrastructure. It was designed for quite the opposite—a diverse set of partner capacities, uneven access to project data, changing reporting requirements, and incomplete evidence at certain project stages. This understanding informed design choice, enabling sustained adoption.

Relational integrity. The insistence that numbers have a traceable story behind them transforms the system from a dashboard into an accountability trail. With a few clicks, a project leader can move from an indicator to the activities that fed it, and from there to the knowledge products generated and linked to evidence. The work becomes easier to document, track, explain, and audit.

Usability at decision-making moments. The system was repeatedly used in ARC funder meetings and internal reviews, not as a redundant reporting artifact, but as a practical interface for managing the conversation. That usability is what drives the maintenance discipline that keeps it alive.

The AI drafting layer: a resource to exploit

One capability that deserves more attention is the use of AI to accelerate narrative reporting. Once Notion holds a resourceful, structured, and evidence-based record, funder-facing narratives can be generated directly from those records. A prompt is built around the technical report template's headings and information structure, and the AI queries the aggregate materials. The result is a first draft already anchored to documented reality, expediting the reporting process.

In the ARC project, this AI drafting layer has not yet replaced the existing reporting format, which is still being used for institutional continuity. But when tested, the results were striking. The report draft accurately mirrored the documented records and dramatically reduced the bottleneck of maneuvering the entire technical team past a blank template. This AI function could be a valuable time-saving resource for researchers when composing funder and technical reports.

The AI output is not the final narrative. It then undergoes group review, and any discrepancy between the draft and the underlying records is clearly visible. The step AI eliminates is the double work of reviewing what has already been reported and wrapping up all new data into coherent language. For a project team, this saves time and is a significant relief.

Challenges and honest limits of the system

The most persistent constraint is evidence readiness. Strong internal documentation is not always immediately shareable. Some evidence sits behind partners' drives, restricted folders, or private links and cannot serve as public-facing proof. The system can be audited internally, while funder transparency requires that the core evidence chain be publicly available. Resolving that tension is inherent to the design work, not an afterthought.

A second challenge is the risk of tool stacking; the perception that teams must enter the same information into multiple institutional systems. The new approach works best when it is clearly positioned as a translation layer that reduces duplication rather than adding yet another layer. In ARC, the metadata structure was designed from the outset to comply with all relevant reporting systems for partners and funders, making interconnectedness a simple step rather than a burden.

The third challenge is the difference between reporting and monitoring incentives. Reporting happens on a schedule; therefore, monitoring needs routines. Without regular cadence, periodic evidence checks, and explicit ownership, the system risks becoming a polished archive: impressive when presented, underused between events. The conditions for failure are clear: momentum collapses when the project leader loses interest, when procedures are not structurally designed and agreed upon across team members, or when incentives and resources are misaligned. The goal is to avoid double work, and when that goal is clearly served, the system sustains itself.

Recommendations

- ✔ **Define a minimal shared standard among project partners.** Establish just enough structure to allow partners to link and compare evidence, without forcing the project into a rigid template. Projects should be free to operate differently, but not free to be untraceable.
- ✔ **Treat public-ready evidence as a design requirement from the start.** When an indicator depends on a private report, the project should plan from the outset how that report will be translated into a shareable brief, a publishable synthesis, or a vetted evidence pack.
- ✔ **Build routines alongside dashboards.** The system delivers value when it is updated in sync with the project life cycle: monthly internal reviews, quarterly evidence checks, and clear roles for who maintains what. The dashboard should become the place the project leader goes to prepare for a funder conversation, not the place someone visits only when asked.
- ✔ **Scale the maintenance resource to the project.** In ARC, a dedicated MEL specialist managed the translation and entry of information. Any team considering this approach should allocate equivalent capacity. The time investment is far less than the time lost to last-minute scrambles for evidence at reporting deadlines.
- ✔ **Exploit the AI drafting layer.** Once the evidence base is structured and linked, AI-assisted narrative generation becomes a genuine productivity tool. Prompt it against the funder template and route the output to group review. The bottleneck this removes is getting a team of scientists past a blank reporting page. Its contribution is real and significant.

Next steps

A practical next step is to pilot the approach in a small set of high-friction, complex contexts where reporting requirements are consistently hard to handle: multi-partner implementation across geographies, requiring frequent funder touchpoints, and blending a mix of technical, socioeconomic, and environmental results.

The pilot should not start by asking teams to adopt a new tool. It should start by defining a minimum evidence chain for each strategic result: Which activity types count? Which deliverables are acceptable proof? Which metadata must always be captured so that it can later be traced?

Once that standard is in place, each project can plug its existing workflows into the same backbone.

For research projects, the value comes from building a clean line, from research design to implementation and dissemination: what was planned, what was done, what outputs were produced, and what is the strongest available evidence that those outputs were taken up or informed decisions. That structure becomes reusable across an entire portfolio, so new projects inherit a consistent, auditable way of packaging evidence without rewriting the logic at each reporting cycle.

The ARC audit now approaching will be the first external validation of the system's integrity. Everything is prepared. That readiness is, in itself, a result.

Reference materials

Paredes D; Ivanova Y. (2025). Agroecological Regenerative Cacao Project Documentation. Notion workspace. <https://bit.ly/42xwAyY>

Bioversity International & CIAT. (2025). Results Dashboard of Outputs, Outcomes, and Impacts [Dashboard]. <https://bit.ly/4uLAbWg>

CGIAR. (2025). Results Dashboard of Windows 1, 2, 3 and Bilateral Portfolio. <https://bit.ly/30XhXly>

Ivanova Y, et al. (2026). Alliance's regenerative agroforestry systems rehabilitating degraded cacao landscapes in the Amazon. Study LAC-2570. Rome (Italy): Bioversity International; Cali (Colombia): International Center for Tropical Agriculture (CIAT). 3 p. <https://hdl.handle.net/10568/183549>

Ivanova Y, et al. (2026). Cacao supply chains in Colombia, Ecuador, and Peru strengthen zero-deforestation commitments through monitoring, traceability, and conservation agreements. Study LAC-2573. Rome (Italy): Bioversity International; Cali (Colombia): International Center for Tropical Agriculture (CIAT). 3 p. <https://hdl.handle.net/10568/183293>

Thomas E, et al. (2026). Unlocking native cacao genetic diversity for resilient and differentiated value chains in Ecuador and Peru. Study LAC-2572. Rome (Italy): Bioversity International; Cali (Colombia): International Center for Tropical Agriculture (CIAT). 3 p. <https://hdl.handle.net/10568/183352>



About Insight Notes

Insight Notes are one of Alliance of Bioversity and CIAT's spaces for sharing learning and results and stimulating conversation and lesson exchanges across programs and projects. The Policy, Inclusion and Socio-Economic Analysis for Impact and Learning (PISA) Program and the Strategic Performance and Results Management (SPRM) Team, jointly with project and program research teams, pull together lessons learned and insights derived from using methodologies, networks, or working along the continuum of impact pathways.

Acknowledgments

This Insight Note was made possible through the financial support of the Fonds Français pour l'Environnement Mondial (FFEM), whose five-year investment in the Agroecological Regenerative Cacao (ARC) project provided both the operational context and the reporting demands that the Results Integration Layer was designed to meet.

The authors thank Olga Spellman of the Alliance of Bioversity International and CIAT Science Editing Unit for her professional editorial inputs and support in reshaping and finalizing this document.

Get in touch

Suggestions and support for developing your own Insight Note: alliance-sprm@groups.cgiar.org

