

RISK → DEAL → HOLD → SUPPLIER

FROM ESG REPORTING TO ESG INTELLIGENCE.

● CONTENTS

What's inside.

Four chapters, one persistent record per entity. The arc follows ESG risk from the document it hides in — before the deal, across the hold, and deep into the supply chain — and shows why reporting is giving way to **intelligence**.

00	Purpose — who this is for	INTRO	01	From the register to risk intelligence	CH 1
02	The register is on a calendar; risk is not	CH 1	03	Three domains, and the fourth that cuts across	CH 1
04	Score the form, or read the document	CH 1	05	The scoring trap	CH 2
06	Two modes of one practice	CH 2	07	The evidence chain builds	CH 2
08	CSDDD: prove it worked	CH 2	09	Four phases, one compounding record	CH 3
10	Documents, not forms	CH 3	11	Six LP-ready reports per investee	CH 3
12	Quant and qual, bound at the record	CH 3	13	Checkbox vs stakeholder intelligence	CH 4
14	Why compliance tools fall short	CH 4	15	Persistent IDs & worker voice at scale	CH 4
16	The through-line	CLOSE	17	See it on your own documents	CLOSE

● PURPOSE · WHO THIS IS FOR

For the teams who answer for ESG.

If you are a risk or ESG lead, an investment partner at an impact or ESG fund, or you run procurement and compliance for a supply chain, this guide is for you. ESG documents arrive faster than any team can read them — and the risk inside them has to be surfaced before it reaches the register, the LP letter, or the regulator.

The argument runs through every chapter: scoring and reporting are no longer the hard part — a register, a rating, a disclosure all draw themselves now. The value moved to the layer that **reads every document on arrival**, scores it against a framework that does not move, and keeps the source sentence behind every flag. That is the shift from ESG reporting to ESG intelligence.

1

ESG risk management

From the annual register to a live reading layer that flags the risk the day it appears.

2

ESG due diligence

Past the scoring trap — an evidence chain on one entity ID, from intake to CSDDD proof.

3

ESG portfolio management

Four compounding phases and six LP-ready reports per investee, generated overnight.

4

Supply chain due diligence

Worker voice read at scale, with persistent IDs that prove a corrective action worked.

i

One record runs the whole arc

The company you screen before the deal is the investee you monitor across the hold and the supplier you audit downstream — one persistent ID per entity, so the evidence compounds instead of resetting at every stage.



CHAPTER ONE

From the ESG register to risk intelligence.

Most ESG risk management is a register reviewed once a year. The risk does not hold still for the review cycle — it appears in a document, on the day that document arrives.

● THE REFRAME

ESG risk management for the risk that moves.

ESG risk management, as most firms practice it, is a register: a matrix for the quarterly committee, a heat map on a slide, a rating bought from an agency. Each is accurate the day it was made — and a little more out of date with every week that follows.

**Scoping a risk and reading for it are two different jobs**

The register scopes the risk, and that scope is sound. But only reading runs fast enough to catch the risk while it can still be managed. ESG risk intelligence is the same discipline run as a live reading layer instead of a periodic register.

ON ARRIVAL

Every document read the day it lands — the founder update, the supplier audit, the diligence file, the board pack.

LOCKED

Scored against an ESG framework that does not move, so March and September are measured by the same ruler.

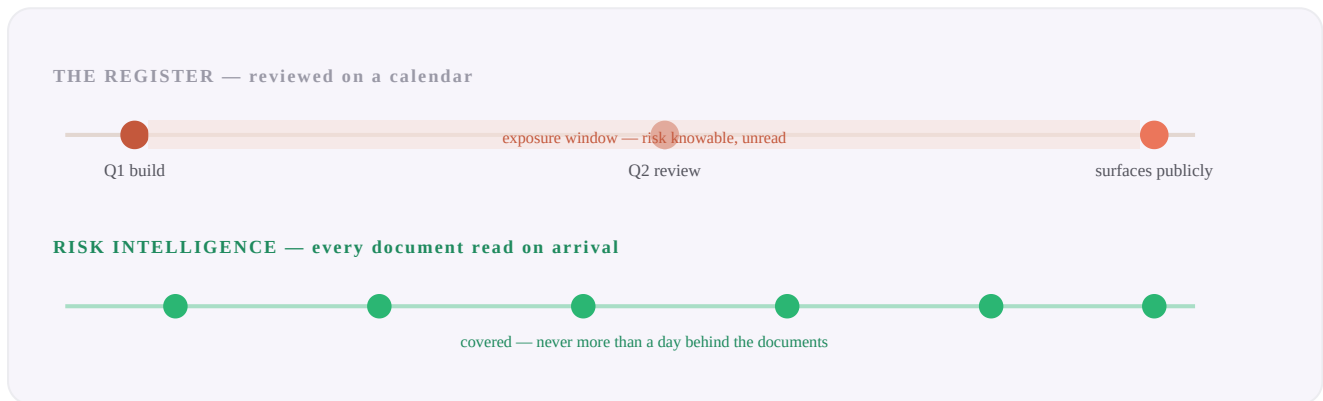
CITED & EARLY

The risk surfaces before the next review — and every flag is traceable to its source sentence.

● THE REGISTER IS REVIEWED ON A CALENDAR. THE RISK IS NOT.

The months between reviews are the months the risk goes **unmanaged**.

It is the same ESG risk on both tracks. The register names it at the next review; risk intelligence names it on arrival.



The register is accurate on the day it is reviewed. Between reviews, it is a record of a risk landscape that has already moved on. **Risk intelligence is a layer, not a calendar.**

● THE RISK TYPES

Three domains of ESG risk — and the **fourth** that cuts across.

Every ESG risk assessment covers the same three domains. What decides whether the assessment works is not which factors are on the list — it is whether the documents behind them were read.

ENVIRONMENTAL

Climate exposure and physical risk · emissions and energy · pollution and waste · resource and water stress · permits · biodiversity.

SOCIAL

Labor practices · health and safety · human rights and modern slavery · diversity and inclusion · community relations · supply-chain labor.

GOVERNANCE

Board structure and independence · anti-bribery and corruption · data privacy and cyber · business ethics · executive pay · disclosure integrity.

4

The fourth risk: evidence

Across all three domains runs a fourth risk — that a claim cannot be evidenced when it matters. An ESG risk you scored but cannot trace to a source is not a managed risk; it is a finding waiting to be questioned. Reading on arrival, with the source kept behind every flag, is how the fourth risk is closed.

● AN ASSESSMENT IS ONLY AS STRONG AS THE READING BEHIND IT

Score the form, or read the document.

An ESG risk assessment has a clean shape — scope, questionnaire, score, report. Most tools run that shape well. They fail at one step inside it: reading. The attachments are where the risk the numbers were chosen to soften tends to sit.

<p>MOST ASSESSMENT TOOLS</p> <p>Scores the form</p>	<p>RISK INTELLIGENCE</p> <p>Reads the document</p>
<p>The questionnaire is scored; the attachments are filed.</p>	<p>Every attachment is read — policy, certificate, audit, free text.</p>
<p>Fast and consistent, but reflects the answer the subject chose to give.</p>	<p>The score is built from the document, not the chosen answer.</p>
<p>A contradiction between policy and audit goes unseen unless an analyst opens the PDF.</p>	<p>Contradictions are flagged, with the source sentence behind each.</p>
<p>Run the same company twice and the rating drifts.</p>	<p>Locked framework — the same question returns the same answer.</p>

! **The one question to ask any AI ESG tool**
 Run the same company twice, a month apart — does the risk rating hold, and can you see the sentence behind it? A locked answer is a finding you can put in the register. A drifting one is a guess with a logo.



CHAPTER TWO


ESG due diligence: past the scoring trap.

Pull a rating on the same company from three providers and the numbers will not agree — sometimes by 25 or 30 points. The fix is not a fourth provider. It is upstream.

● ONE COMPANY, THREE PROVIDERS, ONE MOMENT

A score is fine for screening. Three things break **first**.

Each provider weighs different factors, draws on different disclosures, and updates on a different cycle. As a first-pass anchor, that is fine. As soon as you monitor quarterly, prove diligence to LPs, or stand up CSDDD-grade evidence, the score does less work than it looks like it does.


BREAK 01

The score is already old — most refresh annually. Between pulls, months of news and signal go uncaptured.

BREAK 02

It captures what the company says, not what is happening — ratings lean on the company's own disclosures.

BREAK 03

It cannot prove the diligence worked — a point-in-time number is not a documented chain of action.

● WHAT THE COMPANY SAYS · WHAT IS ACTUALLY HAPPENING

Diligence runs in two modes — on one entity record.

ESG due diligence is the structured assessment of a company’s E, S, and G performance before a commitment, plus the continuous verification of that performance after the relationship begins. The standards are not the hard part — connecting them to the same company across DD, monitoring, and effectiveness proof is.

Mode 01

Intake DDQ

WHAT DOES THE COMPANY SAY IT DOES?

Policies and certifications · 40–150 self-assessment questions · theory of change · closed-ended compliance checks.

Mode 02

Continuous evidence

WHAT IS ACTUALLY HAPPENING ON THE GROUND?

Sustainability documents read end-to-end · worker-voice surveys · audit findings and corrective actions · re-survey vs. the DD baseline.

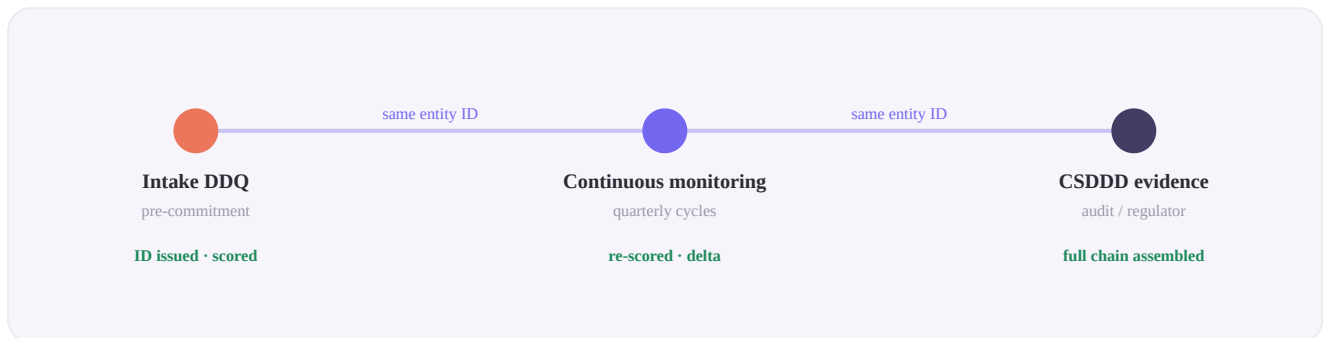

Both modes, one entity record across the lifecycle

One ID per investee or supplier, issued at the first DDQ. Every form, audit, and worker survey after that uses the same ID — so by the CSDDD review, the same row holds the full chain from commitment to verified outcome.

● EVERY STAGE INHERITS THE PRIOR RECORD

One entity ID, issued at the first DDQ, used forever after.

The record gets richer at each stage. By the CSDDD review, the evidence that used to take weeks to reconstruct is one query.



DDQ SCORING

Each score linked to the source passage that supports it — auditable by clicking to the paragraph.

DOCUMENTS READ

Sustainability reports, policy PDFs, certifications, audit findings — read end-to-end against the rubric.

● CSDDD: PROVE IT WORKED

The directive's critical word is “remediate.”

The EU Corporate Sustainability Due Diligence Directive does not just ask companies to identify risks and check boxes. It requires evidence that due diligence is **effective at preventing harm** — and a point-in-time score cannot demonstrate that. Only a documented chain linking commitments at intake to verified outcomes can.

2028

obligations begin for the largest companies (5,000+ staff, €1.5B+ turnover)

3%

of global turnover — the ceiling on penalties, plus civil liability

2027

member-state transposition expected, following the Omnibus I revisions



Assembled by default, not reconstructed at audit

A point-in-time number cannot satisfy the standard. A record of what you saw, when you saw it, what you did, and what changed afterward can — and on one entity ID, that chain is audit-ready by design rather than months of analyst work per cycle.

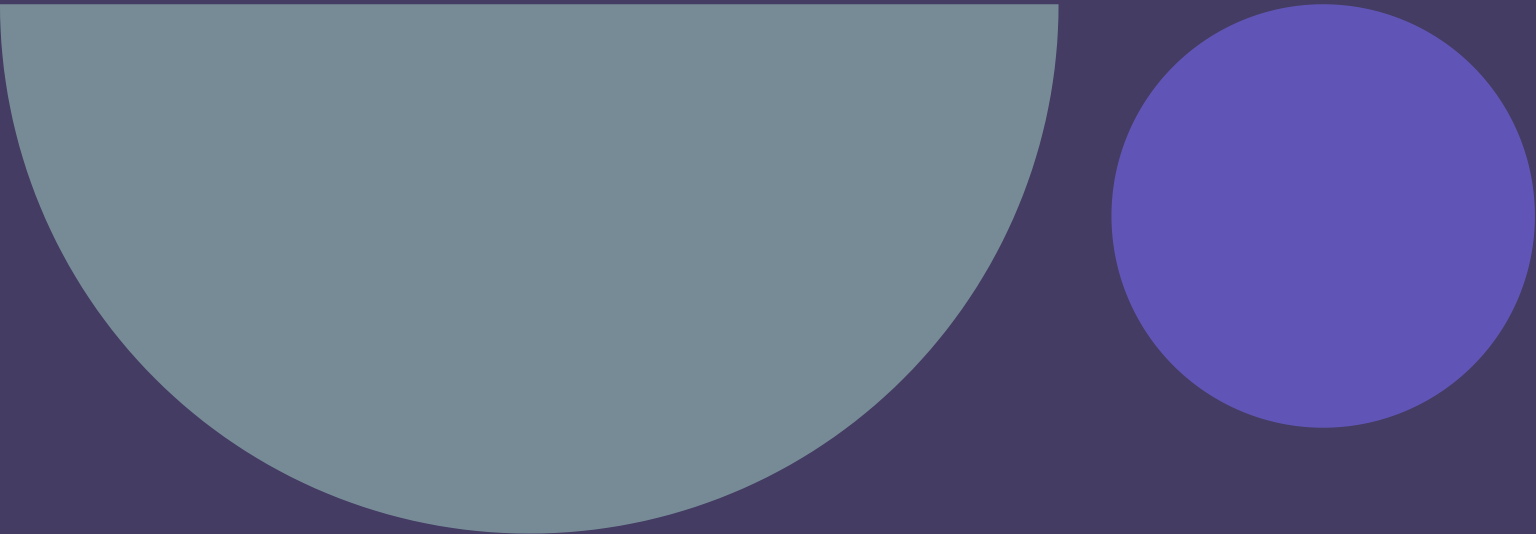
Sopact cites these frameworks — OECD, UN PRI, IRIS+, the Five Dimensions — to share their standard of care, not to certify against them. Compliance is a conversation for your counsel; a defensible, cited ESG record is one this guide can help with.



CHAPTER THREE

ESG portfolio management: four phases, one record.

Investees send documents, not data. LP reports need intelligence, not files. Most funds rebuild it from scratch each quarter — because nothing was indexed last cycle.



● FOUR PHASES. ONE COMPOUNDING RECORD.

Nothing resets at quarter-end. Nothing rebuilds at **exit.**

ESG portfolio work spans the full hold period. Each phase reads more documents, scores more findings, and adds context to the investee record — so the portfolio view at year five is the accumulation of every prior phase, not a snapshot taken at exit.

01 · DD INTELLIGENCE · CONTEXT 5→25%

Pitch deck, financials, theory of change, and founder interviews read at deal close, mapped to your rubric into a queryable investee profile.

02 · LIVING TOC · CONTEXT 25→40%

DD context carries forward. A living theory of change agreed with the investee becomes the reconciliation layer for every quarter.

03 · QUARTERLY LOOP · CONTEXT 40→80%

Investees upload documents in their existing format; Sopact reads them, reconciles against the Living TOC, and generates LP-ready reports overnight.

04 · EXIT & BEYOND · CONTEXT 95%

The investee record is retrievable end to end — first DD document to final exit memo. The exit narrative is a query, not a reconstruction.

● READ WHAT INVESTEEES ALREADY SEND

Documents, not forms.

Quarterly portals get half-filled and miss deadlines. Investees keep sending documents — pitch decks, narrative reports, financials, founder updates. The collection layer should match what investees actually do.

1

One investee record, end to end

Every document, survey, and update ties to the same investee — identity does not break when contacts change or the entity restructures.

2

Compounding context

Every phase inherits the last. Stage resets are why most quarterly cycles burn analyst weeks re-reading the same narrative.

3

The Living TOC is the scoring template

Built with each investee at onboarding; it reconciles every quarter's submission and anchors what "changed" means.

4

Evidence-linked

Every score, trend, and finding traces to a document passage — auditable in two clicks, no black-box scoring.

→

Map to your framework, not a generic one

IRIS+, the Five Dimensions of Impact, SDG alignment, or a custom rubric — a generic ESG scorecard hides what actually matters: outcomes against each investee's specific commitments.

● REPORTING BECOMES A QUERY, NOT A PROJECT

Six LP-ready reports per investee, generated overnight.

If reporting is structured around the investee's actual commitments — read out of DD documents and captured in a Living TOC — the analysis falls out as queries against the accumulated record. The reading work is what most funds underestimate.

01 · SCORECARD

Rubric-mapped scores per investee, each linked to its source.

02 · GAP & RISK MEMO

Narrative anomalies and threshold breaches flagged automatically.

03 · IC BRIEF

One consolidated page per investee — not five open tabs.

04 · LP NARRATIVE

Outcomes against original commitments, reformatted per LP as a query.

05 · LONGITUDINAL TREND

Quarter-over-quarter movement on the same indicators and ID.

06 · EXIT SUMMARY

The hold-period delta, retrievable as one continuous artifact.



Quant and qual, bound at the record

A workforce-diversity drop could be a layoff or a divestiture — only the narrative tells you which. Scores and the narrative that explains each movement live on the same record, read from the same documents, every quarter.



CHAPTER FOUR

Supply chain due diligence: checkbox to intelligence.

Your supplier scored 72 out of 100 on labor practices. But what did workers actually say when they weren't filling in Likert scales? CSDDD will expose the gap.

● COMPLIANCE CHECKBOXES VS STAKEHOLDER INTELLIGENCE

“Is this supplier compliant?” is the wrong question.

Compliance tools answer whether a supplier is compliant. Stakeholder intelligence answers what is actually happening to workers, whether it is getting better, and whether your intervention worked. CSDDD demands the second question.

COMPLIANCE MONITORING Checkbox	STAKEHOLDER INTELLIGENCE Understanding
Structured questionnaire; risk score generated automatically.	Self-assessments, audit PDFs, policies, and worker surveys on one platform, persistent IDs.
Open-ended responses collected but not deeply analyzed.	AI reads policy substance and codes open text across 200+ suppliers at once.
Each assessment is an independent snapshot.	“Forced overtime” up to 12% at three factories — linked to last year’s audit.
Corrective action issued; no proof it changed conditions.	Re-survey shows overtime mentions fell to 2% — the ID proves the connection.

Result, left: “Supplier scored 72/100,” no understanding of why. Result, right: continuous intelligence, evidence that due diligence is effective — **CSDDD-ready**.

● BUILT FOR THE WRONG QUESTION

Three structural gaps in compliance tooling.

The tools that dominate the supply-chain market were built to answer “is this supplier compliant?” — never “what is happening to workers, and are our interventions making it better?”

1

Checkboxes, not understanding

They confirm a code of conduct was submitted. They cannot assess whether a policy is substantive or template language, or read 500 open-ended worker responses for emerging themes.

2

Collected, not analyzed

Worker-voice tools reach millions, but analysis leans on a handful of Likert KPIs. The narrative that explains why a score moved goes unread.

3

No identity across time

Q1 and Q3 are treated as independent snapshots. Proving effectiveness requires connecting the same supplier, factory, and worker cohort across cycles.

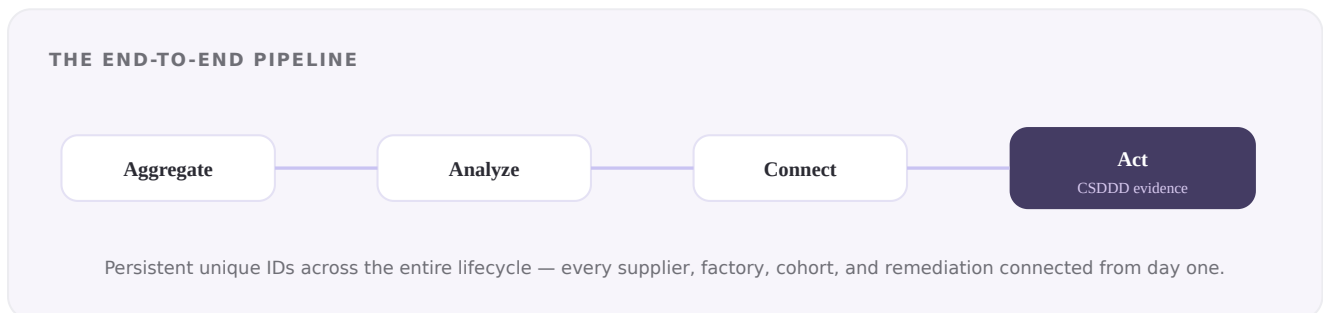
“Due-diligence-oriented technology tools help control risk but rarely identify modern slavery — gaining little trust from workers.”

— BRITISH ACADEMY RESEARCH, ON THE LIMITS OF STRUCTURED SURVEY FORMATS

● PERSISTENT IDS & WORKER VOICE AT SCALE

Proving a corrective action worked.

A persistent ID is assigned to every supplier, factory site, worker cohort, and remediation action from the first interaction — so due diligence, corrective action, follow-up audit, and worker re-survey are all connected. That is the architectural requirement for CSDDD effectiveness proof.



8-12 wk
to produce the report before → minutes after

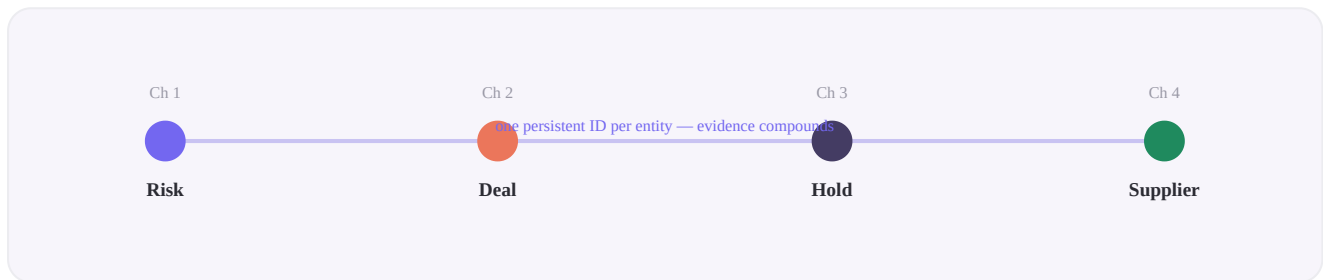
4-5 → 1
data sources unified into one platform

100%
of open-ended responses AI-analyzed, not ignored

● THE THROUGH-LINE

Four chapters, one record.

Risk, deal, hold, and supplier are four stages of the same job: read the documents the subject submits, against one ESG framework, before the risk becomes a loss. The framework, the reading, and the standard of evidence do not change from stage to stage — only the document and the deadline do.

**READ ON ARRIVAL**

Every document — questionnaire, policy, audit, founder update, worker survey — read the day it lands, in any language.

PROVE IT HELD

A locked framework and a cited, longitudinal record — effectiveness an LP, an auditor, or a regulator can check.

● SEE IT ON YOUR OWN ESG DOCUMENTS

Bring the documents. See the risk read on arrival.

Bring a real batch — a company's ESG questionnaire and its attachments, a set of supplier files, or a quarter of investee reports, in whatever languages they arrived. We run it through Sopact and show you the red flags, the contradictions, and every finding traced to its source document.

SOPACT · ESG INTELLIGENCE

Read every ESG document. Score it against one framework. Prove the diligence worked.

A live walkthrough on your real ESG documents — the assessment read on arrival, every flag cited — alongside the risk process you run today. No migration commitment.

[Bring your ESG register →](#)sopact.com[ESG Risk Management](#)[ESG Due Diligence](#)[ESG Portfolio Management](#)[Supply Chain Due Diligence](#)