

AI for ESG and Impact Measurement: Data, Verification and Reporting

An evidence-gated operating model for infrastructure investors and family-office allocators

Prepared by Matchpoint Partners

Author attribution pending CK approval

August 2026

ABSTRACT

Background. Sustainability reporting is becoming digital, taxonomy-based and assurance-ready, while infrastructure investors and family-office allocators still receive fragmented operational, portfolio and impact evidence. **Objective.** This paper develops a controlled AI operating model for A4 infrastructure and digital-infrastructure investors, with A2 family-office CIOs and heads of alternatives as the secondary audience. **Approach.** The analysis reviews 30 primary standards, regulator publications and intergovernmental sources available through 1 August 2026. It connects metric contracts, provenance, digital MRV, deterministic calculation, bounded AI, verification, taxonomy mapping and accountable release. **Findings.** AI can support extraction, classification, reconciliation, anomaly detection, geospatial interpretation and source-linked drafting. Decision-useful disclosure still requires explicit criteria, boundary, source lineage, uncertainty, review and approval. **Implications.** A ninety-day programme should begin with metric identity and shadow operation, then move through representative evaluation, evidence-packet review and bounded release. All worked inputs are unverified illustrative management assumptions; attributed Matchpoint or client revenue, cash cost reduction, loss reduction and alpha remain USD 0 until approved observed evidence exists.

Highlights

Primary user. A4 infrastructure and digital-infrastructure investors; A2 private-capital allocators are the secondary audience.

Evidence unit. A versioned metric packet with criterion, boundary, source, method, controls and approval.

Architecture. Immutable evidence and deterministic calculations sit below bounded AI and reporting adapters.

Verification. Assurance readiness is created upstream through retained provenance, exceptions and segregation of duties.

Economics. Capacity is not cash; revenue, loss and alpha claims require approved observed evidence.

JEL Classification: C45, G23, M41, O32, O33, Q01, Q51

Keywords: artificial intelligence, ESG, impact measurement, infrastructure, family office, sustainability reporting, digital MRV, verification, assurance, provenance, IFRS S1, IFRS S2, ESRS, PCAF, TNFD

Research paper only. This document is not investment, legal, regulatory, accounting, audit, assurance, tax, environmental, engineering, privacy, cybersecurity or technology advice. All worked inputs, thresholds, times, rates and economics are unverified illustrative management assumptions.

1. INTRODUCTION

Sustainability information becomes useful to an investor only when a decision maker can trace a reported figure back to a defined metric, a bounded population, source records, transformations, controls and accountable approval. Artificial intelligence can accelerate parts of that chain. It can extract meter readings and contractual terms, classify evidence, reconcile inconsistent units, identify anomalies, estimate missing values under an approved method, monitor satellite observations and draft source-linked commentary. Each capability creates a new control requirement because the model can also obscure provenance, propagate boundary errors, overstate precision and produce fluent text unsupported by evidence.


This paper develops an evidence-gated operating model for A4 infrastructure and digital-infrastructure investors, with A2 family-office CIOs and heads of alternatives as the secondary audience. The paper addresses physical assets, project companies, funds, direct investments and portfolio reporting. Its primary question is practical: how can an investment team use AI to improve ESG and impact measurement, verification and reporting while preserving the integrity required for investment decisions, lender covenants, regulatory disclosures and independent assurance?

The governing standards supply different decision lenses. IFRS S1 and IFRS S2 focus on sustainability-related risks and opportunities relevant to capital providers [1,2]. GRI focuses on an organisation's impacts on the economy, environment and people [9]. ESRS combines financial materiality and impact materiality, with a digital taxonomy for structured reporting [11-13]. PCAF addresses financed emissions [7]. TNFD addresses nature-related dependencies, impacts, risks and opportunities [10]. ICMA, IFC and IRIS+ provide impact-reporting and impact-management disciplines [16-18]. These lenses can share evidence. Their definitions, scopes, materiality tests and claims must remain distinguishable.

The central proposition is that AI should operate inside an evidence graph rather than becoming the evidence. A metric record should identify the governing criterion, reporting entity, asset boundary, period, unit, source, transformation, uncertainty, reviewer, approval and disclosure destination. Model output should preserve citations to those records. A verifier should be able to reproduce the calculation without relying on a generated narrative.

Research question	Operating answer developed in this paper
Where can AI add value?	In bounded extraction, classification, reconciliation, anomaly detection, forecasting, geospatial interpretation and source-linked drafting.
What must remain authoritative?	Approved criteria, source records, deterministic calculations, review evidence and accountable approval.
How should data quality be represented?	As explicit provenance, coverage, method, uncertainty, freshness and assurance attributes rather than one opaque score.
How should impact claims be controlled?	Separate intent, output, outcome, contribution, additionality and risk; require evidence for every transition.
How should productivity be measured?	Accepted evidence packets per paid hour, including review, rework, exceptions and assurance preparation.
When may financial value be attributed?	After an approved baseline and observed evidence establish the causal bridge. Until then, attributed revenue, cash cost reduction, loss reduction and alpha remain USD 0.

The analysis reviews 30 primary standards, official regulator publications and intergovernmental sources available through 1 August 2026. It develops an original Matchpoint architecture, metric contract, control catalogue, worked scenarios and ninety-day adoption roadmap. The worked scenarios are explicitly unverified illustrative management assumptions. They do not represent Matchpoint clients, realised savings, investment performance or assured disclosures.

ESG evidence-to-assurance chain**Figure 1. ESG evidence-to-assurance chain**

Source: Matchpoint framework; standards and assurance context [1-9,30].

2. THE A4 AND A2 DECISION PERIMETER

2.1 A4 infrastructure investors

A4 investors allocate capital to energy, transport, utilities, data centres and other real-asset infrastructure. Their sustainability information sits close to physical operations. Electricity meters, fuel records, water systems, construction schedules, environmental permits, grid factors, equipment specifications, tenant activity and geospatial observations may all contribute to a reported metric. The same asset can support several decisions: acquisition underwriting, project finance, covenant monitoring, operating improvement, valuation, insurance, exit preparation and external reporting.

The data burden is therefore multi-level. A project company needs asset-level controls. A fund manager needs consistent portfolio aggregation. A limited partner needs comparable, decision-useful information. A lender needs covenant definitions and evidence. An assurance practitioner needs sufficient appropriate evidence against suitable criteria [8,30]. A single generated dashboard cannot satisfy these roles unless its lineage and decision rights are explicit.

2.2 A2 family-office allocators

A2 allocators often receive heterogeneous reports from managers, direct holdings and advisers. They may see the same concept expressed through incompatible boundaries, periods, units and estimation methods. One manager reports location-based electricity emissions, another reports market-based emissions, a third provides total carbon dioxide equivalent without a methods note. A direct asset reports avoided emissions against a proprietary counterfactual. A portfolio company provides a narrative about community benefit without baseline or outcome evidence.

AI can help an A2 team normalise these packages, map disclosed metrics to a house ontology, identify missing evidence, compare methods and prepare investment-committee questions. The team should preserve the manager's original representation and label every allocator adjustment. An allocator normalisation is an analytical view. It does not become the investee's reported fact or an assured measure.

2.3 Decision rights

Decision layer	A4 asset or fund role	A2 allocator role	Required human authority
Criterion	selects applicable standard, covenant or methodology	defines house information requirements	policy owner, legal or reporting lead
Boundary	establishes entity, asset, period and value-chain scope	records manager scope and allocator normalisation	asset controller, fund controller or CIO delegate
Measurement	captures source observations and deterministic calculations	ingests manager and direct-asset evidence	data owner and metric owner
AI processing	extracts, maps, flags and drafts	normalises, compares and prepares questions	named workflow owner

Decision layer	A4 asset or fund role	A2 allocator role	Required human authority
Verification	tests evidence, method and controls	challenges manager evidence and house adjustments	independent reviewer or assurance practitioner
Disclosure	releases external statement or investor report	approves committee and beneficiary reporting	accountable executive, board or authorised committee

AI may recommend a classification or identify an anomaly. It should not approve a reporting boundary, certify a climate claim, determine materiality, sign an assurance conclusion or release a public disclosure. Those acts carry institutional accountability.

A4 and A2 sustainability decision map




Figure 2. A4 and A2 sustainability decision map

Source: Matchpoint ICP and decision-rights framework.

3. STANDARDS, MATERIALITY AND CLAIM TYPES

3.1 Map the decision lens before the metric

IFRS S1 and IFRS S2 establish an investor-oriented baseline for sustainability-related financial information [1,2]. Their structure covers governance, strategy, risk management, and metrics and targets. The IFRS Sustainability Disclosure Taxonomy supports machine-readable tagging and is designed to work with the IFRS Accounting Taxonomy [3]. Taxonomy updates triggered by amended greenhouse-gas disclosures illustrate the need for version-controlled mappings [4].

GRI requires reporting organisations to identify material impacts and apply principles including accuracy, balance, clarity, comparability, completeness, timeliness and verifiability [9]. ESRS uses double materiality and a detailed digital taxonomy [11]. The European Commission adopted revised ESRS on 3 July 2026; the revised package reduces mandatory datapoints and adds a voluntary standard for smaller undertakings, subject to the stated EU scrutiny process [12,13]. A production system must therefore record the applicable reporting version and effective period.

PCAF provides methods for attributing financed emissions to financial activities and distinguishes core inventories from optional forward-looking or avoided-emissions information [7]. TNFD organises nature-related disclosure around governance, strategy, risk and impact management, and metrics and targets [10]. ICMA recommends transparent use-of-proceeds and impact reporting for green bonds [16]. IFC's Operating Principles connect strategic intent, investment selection, monitoring, exit and independent verification [17]. IRIS+ provides a practical metric-selection system around what, who, how much, contribution and risk [18].

These frameworks do not form one interchangeable metric library. They create a standards map.

Claim class	Typical question	Governing evidence	AI boundary
Financial risk	Could the issue affect cash flows, access to finance or cost of capital?	financial model, risk register, asset exposure, scenario assumptions	extract and link evidence; no autonomous materiality conclusion

Claim class	Typical question	Governing evidence	AI boundary
Operational footprint	What emissions, energy, water or waste arose within the boundary?	meter, invoice, activity data, emission factor, calculation	reconcile and flag; deterministic calculation remains reproducible
Financed emissions	What share of emissions is attributed to financing or investment?	PCAF method, outstanding amount, enterprise value or project denominator, emissions data	map and calculate under approved method; disclose data quality
Output	What activity or product was delivered?	operational record, beneficiary count, capacity or service unit	extract and deduplicate under approved definitions
Outcome	What changed for people or planet?	baseline, follow-up measurement, counterfactual context, stakeholder evidence	support analysis; uncertainty and alternative explanations remain visible
Contribution	What part of the change is attributable to the enterprise or investor?	causal design, additionality case, contribution analysis	assist modelling; no generated causal assertion
Public narrative	What can be stated externally?	approved claims register and cited evidence packet	draft from approved facts; human release required

3.2 The claims ladder

A strong system prevents a low-level observation from being promoted into a high-level claim without evidence. A smart meter can establish an electricity reading within its calibration and coverage limits. An emission factor can convert activity data to an emissions estimate under a specified method. A year-on-year reduction can be calculated after boundary and method consistency are checked. A claim that management action caused the reduction requires further evidence. A claim that an investment generated environmental impact requires a contribution analysis and impact-risk assessment.

The claims ladder should be encoded as permissible transitions:

1. source observation;
2. validated activity datum;
3. calculated metric;
4. period comparison;
5. target-progress assessment;
6. risk or financial interpretation;
7. outcome statement;
8. contribution or additionality claim;
9. externally released disclosure.

Every transition needs an owner, method and evidence threshold. AI may automate parts of transitions one to five under controls. Transitions six to nine require substantive judgement and approval.

4. THE ESG METRIC CONTRACT

4.1 One record for each metric version

The metric contract is the core operating record. It separates the label displayed in a report from the definition that makes the number meaningful.

Field	Required content
Metric ID	stable identifier plus version
Claim class	footprint, financed emissions, output, outcome, contribution, risk or narrative
Governing criterion	standard, regulation, covenant, policy or methodology and version
Reporting entity	legal entity, fund, portfolio, project or asset
Boundary	organisational, operational, geographic, value-chain and beneficiary boundary
Period	measurement start, end, cut-off and restatement status
Unit	base unit, display unit, currency, conversion and rounding rule
Source	system, document, meter, survey, satellite product or external dataset

Field	Required content
Transformation	extraction, mapping, conversion, aggregation and calculation steps
Estimation	method, reason, population affected and replacement plan
Uncertainty	measurement, model, sampling and scenario uncertainty where relevant
Data quality	coverage, freshness, completeness, provenance and assurance attributes
Control	validation, reconciliation, exception and approval evidence
Owner	source owner, metric owner, reviewer and release authority
Disclosure mapping	IFRS, ESRS, GRI, PCAF, TNFD, ICMA, lender or internal destination
Retention	source and evidence-pack retention requirement

4.2 Metric identity prevents silent drift

Metric identity should change when a material definition changes. A new emission factor, revised asset boundary, updated PCAF method, new beneficiary definition or model version should create a new metric version. Restatement policy should specify whether historical periods are recalculated and how both versions remain discoverable.

A report label such as "renewable energy generated" is insufficient. The contract should state gross or net generation, meter location, auxiliary consumption, curtailment treatment, export treatment, reporting period, data gaps and whether any estimate was used. AI mapping should return the metric ID and confidence with the extracted value. A value without a valid metric identity should enter an exception queue.

4.3 Data quality is a vector

One composite quality score can conceal the reason a number is weak. The system should preserve a vector:

- provenance: original, copied, inferred or externally sourced;
- coverage: share of assets, time or population represented;
- completeness: required fields present;
- freshness: lag between event, capture and use;
- method: direct measurement, calculation, estimation or proxy;
- control: validation and reconciliation status;
- uncertainty: known measurement or model range;
- assurance: unassured, internally reviewed, limited assurance or reasonable assurance where applicable;
- comparability: boundary and method consistency across periods or assets.

PCAF's data-quality concepts show why measured and estimated financed-emissions data should remain distinguishable [7]. GRI's reporting principles likewise support explicit accuracy, completeness and verifiability attributes [9]. A2 allocators can use the vector to direct engagement: improve coverage at one manager, replace estimates at another, and resolve boundary inconsistency at a third.




Figure 3. Metric provenance and claim graph

Source: Matchpoint evidence-graph framework; provenance context [26].

5. SOURCE DATA AND DIGITAL MEASUREMENT

5.1 Source hierarchy

The source hierarchy should prefer records closest to the event and preserve corroborating observations. Examples include:

Evidence tier	Examples	Control focus
Primary operational	calibrated meter, invoice, fuel record, payroll record, safety system	completeness, calibration, cut-off, ownership
System-derived	ERP ledger, building management system, data-centre telemetry, project-control system	access, interface, master data, change log
Document-derived	environmental permit, power-purchase agreement, contractor report, policy	version, signature, effective date, extraction evidence
Geospatial or sensor	satellite plume, land-cover classification, remote temperature, IoT reading	resolution, coverage, attribution, cloud or noise effects, model version
Survey or stakeholder	beneficiary survey, employee response, community consultation	sampling, consent, question design, non-response and representativeness
External factor	grid factor, emission factor, sector benchmark, taxonomy	issuer, jurisdiction, vintage, applicability and update
Modelled estimate	proxy, interpolation, forecast or counterfactual	documented method, uncertainty, approval and replacement plan

5.2 Digital MRV

The World Bank defines digital monitoring, reporting and verification as the use of digital technologies and processes to automate data collection, emissions reporting and the generation of mitigation outcomes [19,27]. Its 2025 guidance addresses system evaluation, interoperability and control hotspots. Digital MRV can connect meters, sensors, satellites, mobile applications, cloud systems and registries. Reporting uses should preserve the records required by the applicable transparency arrangements [21]. The architecture gains value when the physical and digital layers are bound by consistent identities, controls and audit trails.

UNEP's Methane Alert and Response System provides an operational example. As of July 2026, MARS uses data from more than 30 satellite instruments, scientific expertise and advanced AI models to detect and notify governments and companies of very large methane emissions; its scope expanded to coal and waste in 2026 [20,28]. The example demonstrates scalable AI-assisted observation. It also shows the importance of method documentation, source attribution, response workflow and corroboration. A satellite alert is a decision signal whose asset attribution and mitigation response require operational follow-up.

5.3 The physical-digital identity bridge

Every asset, meter, project, legal entity, financing instrument and reporting boundary should have a stable identifier. The identity bridge connects:

- legal entity to fund and ownership period;
- asset to location and operating boundary;
- meter to equipment, unit and calibration record;
- document to issuer, signature and effective period;
- investment to financing amount and PCAF attribution denominator;
- beneficiary record to consent, programme and deduplication rules;
- geospatial observation to coordinates, time, resolution and confidence.

AI entity resolution can propose links. High-impact or ambiguous links should require review. The evidence graph should retain rejected alternatives and the final approver. W3C PROV supplies a useful generic model of entities, activities and agents for interoperable provenance [26].

6. THE CONTROLLED AI ARCHITECTURE

6.1 Seven layers

The target architecture has seven layers:

1. source and identity;
2. immutable raw evidence;
3. validated and normalised data;
4. deterministic metric calculation;
5. bounded AI services;
6. review, approval and assurance workspace;
7. reporting, taxonomy tagging and disclosure.

The raw layer preserves original files, observations and metadata. The validated layer applies units, identities, cut-off and quality rules. Deterministic calculations implement approved formulas. AI services can then extract, classify, match, detect anomalies, forecast or draft from controlled inputs. Reviewers see the generated output beside cited evidence, method and exceptions. Reporting adapters map approved metric versions to the relevant taxonomy or template.

Controlled ESG and impact data architecture



Figure 4. Controlled ESG and impact data architecture

Source: Matchpoint architecture; digital MRV and AI-governance context [19,22,23,27].

6.2 AI service catalogue

AI service	Permitted task	Required evaluation	Release boundary
Document extraction	extract dates, units, values, clauses and evidence spans	field precision and recall by document type; citation fidelity	low-confidence and material fields reviewed
Classification	map source or metric to ontology	class-level error, abstention and drift	novel or ambiguous class enters exception queue
Entity resolution	link asset, issuer, project, meter or beneficiary	false-link rate by risk tier	material links require approval
Reconciliation	propose matches and explain breaks	accepted-match rate, false match, unresolved population	deterministic tolerance and reviewer sign-off
Anomaly detection	flag unusual readings, gaps or relationships	detection rate, false-alert burden, stability	signal only; no automatic restatement
Geospatial analysis	identify land cover, plume or physical change	spatial resolution, ground truth, class error and coverage	corroboration and domain review required
Forecasting	estimate target path, physical risk or missing period	backtest, calibration, interval coverage and drift	scenario or estimate label preserved
Generative drafting	prepare source-linked report text or questions	citation support, omission, contradiction and prohibited-claim tests	approved facts only; human release

NIST's AI RMF organises risk management around govern, map, measure and manage [22]. Its generative-AI profile identifies information-integrity and provenance risks relevant to public reporting [23]. The architecture should therefore inventory models, record intended use, test representative data, monitor change and keep incident, override and decommissioning processes.

6.3 Retrieval and citation

Retrieval should operate on approved evidence packets rather than an unrestricted document universe. Each chunk should retain document ID, version, page or location, effective date, security label and metric relationship. Generated text should cite the exact evidence span or metric record. Unsupported sentences should be blocked or returned for manual completion.

The system should distinguish three states:

- supported: the statement is entailed by approved evidence;
- qualified: evidence supports the statement within a stated method, boundary or uncertainty;
- unsupported: evidence is absent, contradictory or below the required threshold.

The drafting interface should favour abstention over fluent completion. External report text should be generated only from supported or approved qualified statements.

7. VERIFICATION AND ASSURANCE READINESS

7.1 Verification is designed upstream

Assurance cannot repair missing provenance after publication. The evidence pack should be created with the metric. It should include the governing criterion, source population, extraction result, rejected records, transformation log, formula, factor version, reconciliation, exception resolution, reviewer sign-off and final disclosure mapping.

ISSA 5000 applies across sustainability topics and suitable reporting frameworks, and supports limited and reasonable assurance engagements [8,30]. The production system should not label a metric "assured" unless the assurance conclusion covers that information and period. Internal review, model validation and independent assurance are different control layers.

7.2 Evidence sufficiency matrix

Claim	Minimum internal evidence	Additional assurance focus
Measured activity	complete source population, calibration or system control, cut-off	source reliability and completeness
Calculated footprint	approved method, factors, units, boundary and reproducible calculation	criteria suitability and transformation accuracy

Claim	Minimum internal evidence	Additional assurance focus
Estimated metric	rationale, method, input quality, uncertainty and replacement plan	estimation bias and disclosure adequacy
Financed emissions	investment exposure, denominator, emissions data and PCAF method	attribution method, data-quality disclosure and consistency
Outcome	baseline, follow-up, population, method and adverse effects	measurement design and alternative explanations
Contribution	counterfactual or contribution reasoning, evidence and risk	causal support and risk of overstatement
Narrative	approved claims register and evidence citations	consistency with underlying facts and balanced presentation

7.3 Segregation of duties

The source owner should not unilaterally approve the metric, alter the model and release the disclosure. At minimum, the system should separate:

- source submission;
- metric preparation;
- model or rule administration;
- review and exception approval;
- external release;
- independent assurance access.

Privileged changes require a ticket, reason, test evidence and effective date. The evidence graph should show which version produced each disclosure. A post-period model update should not silently rewrite the historical evidence pack.

8. EMISSIONS, ENERGY AND RESOURCE MEASUREMENT

8.1 Emissions inventory

GHG Protocol provides corporate and value-chain accounting foundations [5,6]. The workflow should establish organisational and operational boundaries before loading data. It should record Scope 1, Scope 2 and relevant Scope 3 categories; location-based and market-based Scope 2 views should remain distinct when required. Activity data, factor source, factor vintage, global-warming potential basis, unit conversion and estimation status should be traceable.

AI can extract consumption from invoices, map facilities to grid factors, identify missing periods and reconcile energy records to the general ledger or operations systems. Deterministic rules should perform the emissions calculation. A model-generated factor or undocumented web value is unacceptable. Factor updates should create a controlled restatement decision.

8.2 Infrastructure operating metrics

The same design applies to energy generated, energy consumed, water withdrawn, water discharged, waste, refrigerants and data-centre efficiency metrics. Asset context matters. A data centre should connect utility invoices, on-site generation, uninterruptible power systems, cooling, tenant loads and water systems to defined boundaries. A renewable asset should distinguish gross generation, auxiliary load, curtailment and exported energy.

An anomaly engine can flag a discontinuity. Operations must investigate whether it represents equipment performance, meter error, missing data, boundary change or ordinary seasonality. The metric record should preserve the investigation and resolution.

8.3 Financed emissions

PCAF Part A supplies methods for financed emissions across financial asset classes, including project finance and commercial real estate [7]. An A4 fund should link the exposure amount, attribution denominator, asset emissions and data-quality attributes for each reporting date. An A2 allocator should preserve whether a manager supplied reported emissions, an estimate or a portfolio-level allocation.

Portfolio aggregation should prevent double counting across ownership structures and periods. If a fund owns a project through multiple vehicles, the legal and economic ownership graph should resolve the exposure once under the selected method. AI can assist entity resolution and classification; the approved ownership and attribution rules remain deterministic.

9. NATURE, SOCIAL AND IMPACT MEASUREMENT

9.1 Nature-related evidence

TNFD recommends disclosure across governance, strategy, risk and impact management, and metrics and targets [10]. Infrastructure projects can affect land, water, ecosystems and communities at specific locations. Geospatial data and machine learning can support location screening, land-cover monitoring and change detection. Their evidence pack should disclose imagery source, date, resolution, classification model, ground truth, cloud or obstruction treatment, uncertainty and reviewer interpretation.

A detected land-cover change does not by itself establish causation, legal breach or material financial effect. The system should connect the observation to permits, project boundaries, field evidence and responsible review. Nature dependencies and impacts may require domain expertise that a general model cannot supply.

9.2 Social and beneficiary information

Social and impact measures may involve workers, contractors, customers and affected communities. The metric contract should address privacy, consent, safeguarding, sampling, representativeness and grievance mechanisms. Deduplication should not erase meaningful repeat participation. Sensitive attributes should be access-controlled and used only for approved purposes.

AI can code survey responses, translate text and identify themes. The research team should test language, subgroup and cultural performance. Generated summaries should include minority and adverse views where material. A balanced report cannot rely on sentiment aggregation alone.

9.3 Impact logic

IFC's Operating Principles and IRIS+ support a structured chain from strategic intent through measurement and management [17,18]. The evidence graph should represent:

- input: capital, labour or resource committed;
- activity: work performed;
- output: product, service or capacity delivered;
- outcome: change experienced by people or planet;
- contribution: the enterprise or investor's role in that change;
- impact risk: the chance that outcomes differ from expectations or create harm.

The system should record negative and unintended effects. Impact reporting should not be limited to selected positive indicators. AI can monitor the chain and identify inconsistent claims; impact ownership and remedial action remain human responsibilities.




Figure 5. AI use-case and control matrix

Source: Matchpoint use-case and decision-rights framework.

10. REPORTING, TAXONOMY AND DISCLOSURE

10.1 One evidence base, controlled reporting views

The reporting layer should render approved metric versions into framework-specific views. It should preserve the relationship between a source fact and each disclosure mapping. This allows one electricity-consumption record to support a GHG calculation, an operational KPI, an investor dashboard and a disclosure tag while keeping the relevant methods and claim types distinct.

The system should never infer equivalence solely from similar labels. A mapping record should state:

- source metric ID and version;
- target standard, paragraph, datapoint or taxonomy element;
- transformation or aggregation;
- dimensional members such as geography, asset, period or Scope 3 category;
- mapping owner and approval;
- effective date and superseded mapping;
- disclosure text or table destination.

IFRS and ESRS digital taxonomies demonstrate how tags and dimensions can support extraction and comparison [3,11]. IFRS Foundation supporting materials provide implementation resources for applying the ISSB Standards [29]. A reporting engine should validate required dimensions, sign, unit, period type and duplicate facts before release. Taxonomy validation complements substantive review. A technically valid XBRL fact can still be wrong, incomplete or unsupported.

10.2 Narrative controls

Narrative drafting carries a distinctive risk because a plausible sentence can imply more than the metrics establish. The claims register should contain the approved subject, predicate, qualifiers, evidence IDs, permitted reporting destinations and expiry or review date. Generated commentary should use only registered claims and current approved metrics.

Examples of mandatory qualifiers include:

- estimated or measured;
- reported by investee or independently verified;
- gross or net;
- location-based or market-based;
- operational footprint or financed emissions;
- output or outcome;
- expected or realised;
- portfolio coverage percentage;
- material exclusions;
- uncertainty or restatement.

The European Commission's July 2026 AI Act guidance addresses transparency for certain AI-generated public-interest content from 2 August 2026 [24]. Applicability depends on the system, role and use. Legal review remains necessary. The operating rule is simpler: record AI involvement in public reporting, preserve human editorial approval and maintain the evidence supporting every material statement.

10.3 UAE reporting context

The UAE Sustainable Finance Working Group's disclosure principles call for adequate systems, procedures and data governance to monitor sustainability matters and address data gaps [14]. The CBUAE climate-risk regulation requires financial institutions to develop governance, data, measurement and reporting capabilities for climate-related financial risks [15]. An A4 manager or A2 allocator working with regulated UAE institutions should expect evidence requests to extend beyond a polished annual report into data quality, methodology, scenario analysis and governance.

The reporting engine should therefore support both external disclosures and controlled lender or due-diligence evidence rooms. Sensitive asset data may be shared through role-based access, redacted views and logged downloads. Public transparency and information security are simultaneous design requirements.

Before and after: ESG reporting workflow**Figure 6. Before and after: ESG reporting workflow**

Source: Matchpoint operating-model framework.

11. THE BEFORE-AND-AFTER OPERATING MODEL**11.1 Fragmented current state**

A common current-state pattern has local spreadsheets, emailed templates, manual document review, inconsistent metric labels and narrative preparation late in the reporting cycle. Reviewers ask for evidence after aggregation. Version history sits in filenames. Restatements are difficult to reproduce. Portfolio teams spend time finding and rekeying data, while assurance teams reconstruct lineage.

This state creates four forms of hidden work:

1. acquisition work: chasing files and resolving permissions;
2. semantic work: understanding definitions, boundaries and units;
3. control work: identifying gaps, duplicates and inconsistencies;
4. presentation work: translating evidence into multiple reporting formats.

Hours alone do not establish value. Some activity is required judgement. Some is avoidable rework. Some should be performed by a domain expert. The baseline should classify the work before automation.

11.2 Controlled target state

The target state begins with an approved metric catalogue and submission contract. Sources enter through controlled connectors or evidence uploads. Automated validation checks identity, period, unit, completeness and expected relationship. AI extraction and classification operate on bounded document types. Exceptions route to the accountable owner. Deterministic calculations produce versioned metrics. Reviewers approve evidence packets. Reporting adapters create tagged tables and source-linked narrative drafts.

Workflow stage	Current-state evidence	Target-state control	Acceptance measure
Request	emailed template and reminders	metric-specific submission contract	submissions received by cut-off
Capture	copied values and attachments	immutable source plus metadata	source population reconciled
Understand	analyst interprets label	approved ontology and mapping	correct metric ID and boundary
Validate	manual spot check	deterministic rules plus exception queue	exceptions resolved or disclosed
Calculate	local spreadsheet	versioned deterministic engine	reproducible result and factor set
Review	comments across files	evidence packet and structured sign-off	review complete with audit trail
Report	copied tables and narrative	framework adapters and cited drafting	approved facts, tags and qualifiers

Workflow stage	Current-state evidence	Target-state control	Acceptance measure
Assure	evidence assembled on request	retained assurance workspace	requested evidence available and consistent

11.3 Human work in the target model

Human effort shifts toward method, exceptions, evidence challenge, domain interpretation and decision-making. The operating model should not count a generated draft as accepted output. A metric packet is accepted only when required sources, validation, review and disclosure mapping pass. A narrative is accepted only when every material statement is supported, qualified and approved.

12. TWO UNVERIFIED ILLUSTRATIVE SCENARIOS

Every input in this section is an unverified illustrative management assumption. The examples do not describe Matchpoint clients or observed engagements. They demonstrate calculation and control logic. Attributed Matchpoint or client revenue, cash cost reduction, loss reduction and alpha remain USD 0.

12.1 Scenario A: A4 digital-infrastructure fund

[Unverified illustrative scenario] A fund owns eight digital-infrastructure assets. The reporting team receives utility invoices, meter exports, tenant schedules, environmental permits and quarterly operating packs. The team prepares energy, water, Scope 1, Scope 2 and selected Scope 3 metrics, plus lender and investor commentary.

Input	Unverified assumption
Assets	8
Reporting cycles per year	4
Metric packets per asset per cycle	28
Annual metric packets	896
Baseline preparation and review hours	1,920
Loaded planning rate	USD 75 per hour
One-time programme cost	USD 210,000
Annual run cost	USD 96,000

The pilot limits AI to invoice and document extraction, unit mapping, missing-period detection and source-linked narrative drafting. Emissions calculations, factor selection, boundary approval and external release remain deterministic or human-controlled.

[Unverified calculation] The team models 1,056 target-state hours, comprising 720 preparation hours, 240 review hours and 96 exception or administration hours. Modelled gross capacity released is:

1,920 baseline hours - 1,056 target-state hours = 864 hours

[Unverified calculation] Planning capacity value is:

864 hours x USD 75 per hour = USD 64,800

This amount is capacity value, not cash cost reduction. The scenario provides no approved evidence of headcount removal, avoided external spend or revenue contribution. Attributed cash cost reduction is therefore USD 0. A decision to proceed would need control, reporting-quality, assurance-readiness and strategic benefits in addition to any observed capacity value.

Pilot gate	Unverified target	Evidence required
Material-field extraction precision	at least 98%	labelled representative test set
Evidence citation fidelity	100% for accepted packets	exact source-location check
False entity link	at most 0.2%	asset and meter identity test
Exception closure	at least 95% before report lock	signed exception register
Metric reproducibility	100%	independent rerun from source and method

Pilot gate	Unverified target	Evidence required
Unsupported external sentences	0	claims-register and citation test

These thresholds are unverified management proposals. The accountable owner should approve or replace them based on materiality, reporting use and risk tolerance.

12.2 Scenario B: A2 family-office portfolio

[Unverified illustrative scenario] A family office receives quarterly ESG and impact packages from eighteen managers and six direct holdings. Each package uses different templates. The investment team wants a house view of coverage, financed emissions, energy exposure, impact indicators and controversies without altering the underlying manager representations.

Input	Unverified assumption
External managers	18
Direct holdings	6
Reporting cycles per year	4
Submitted metric rows per cycle	2,400
Annual metric rows	9,600
Baseline normalisation and review hours	1,440
Loaded planning rate	USD 90 per hour
One-time programme cost	USD 165,000
Annual run cost	USD 78,000

The target workflow retains every original value, manager label, method note and file. AI proposes house-ontology mappings and flags absent boundary, period, unit, methodology, coverage or assurance fields. A reviewer approves the mapping. Portfolio calculations use approved mappings only. The committee dashboard displays both manager-reported and allocator-normalised views.

[Unverified calculation] The team models 760 target-state hours, comprising 420 preparation hours, 260 review hours and 80 exception or administration hours. Modelled capacity released is:

$$1,440 \text{ baseline hours} - 760 \text{ target-state hours} = 680 \text{ hours}$$

[Unverified calculation] Planning capacity value is:

$$680 \text{ hours} \times \text{USD } 90 \text{ per hour} = \text{USD } 61,200$$

The scenario provides no evidence that capacity is converted to cash savings or superior investment performance. Attributed cash cost reduction, loss reduction and alpha are USD 0. The decision case should include observed improvements in evidence coverage, question quality, decision cycle and control reliability without converting them into financial value until approved causal evidence exists.

House view	Required label
manager value used without change	Manager reported; unassured unless assurance evidence states otherwise
allocator unit conversion	Allocator normalisation; method shown
allocator estimate	Allocator estimate; input and uncertainty shown
AI-proposed mapping	Pending until reviewed; excluded from approved portfolio totals
missing data	Missing; no zero substitution
conflicting manager versions	Conflict open; both versions retained

12.3 Scenario boundary

Both scenarios demonstrate that productivity economics can be calculated before value is realised. A planning rate converts time to capacity value. Cash, revenue, loss and alpha attribution require further observed evidence. The programme should retain zero for those fields until the evidence gate is met.

13. PRODUCTIVITY, QUALITY AND VALUE ATTRIBUTION

13.1 Accepted evidence packets

The primary productivity unit is an accepted evidence packet. It includes source records, metric identity, validation, deterministic calculation where required, AI output, exception resolution, reviewer approval and disclosure mapping. Counting extracted fields, generated sentences or model calls rewards activity rather than decision-useful output.

Use these measures:

Accepted-packet productivity = accepted packets / total paid hours

First-pass acceptance = packets accepted without material rework / packets reviewed

Exception burden = exception handling hours / total workflow hours

Evidence retrieval time = reviewer minutes from request to sufficient evidence


Unsupported-claim rate = unsupported material claims / material claims reviewed

Coverage = represented population / required population

13.2 Value ledger

Value field	Evidence gate	Initial attributed value
capacity	accepted baseline and target workflow time	observed minutes only
cash cost reduction	approved avoided payroll, contractor or assurance spend	USD 0
revenue	observed incremental pricing, conversion, retention or service capacity with approved attribution	USD 0
loss reduction	observed avoided loss against approved counterfactual method	USD 0
alpha	approved performance attribution net of fees, risk and benchmark	USD 0
reporting-quality improvement	approved quality metric and baseline	observed non-financial measure only
assurance-readiness improvement	approved evidence-retrieval or exception metric	observed non-financial measure only

The ledger should record baseline window, population, acceptance rule, data exclusions, implementation cost, run cost, reviewer effort and confidence. Capacity should remain separate from cash. Better coverage should remain separate from performance. An investment decision informed by better evidence does not establish alpha.

Evidence-gated productivity and value bridge**Figure 7. Evidence-gated productivity and value bridge**

Source: Matchpoint measurement and attribution framework.

13.3 Release gates

1. purpose and reporting criteria approved;
2. metric catalogue and boundary approved;
3. representative source population tested;
4. identity, access and retention controls approved;
5. AI evaluation passes by use case and risk tier;
6. deterministic calculations reproduce independently;
7. evidence citations and exceptions pass review;
8. assurance-readiness review completed where relevant;
9. value ledger records observed results and zeroes unsupported fields;
10. accountable authority approves bounded production use.

14. FAILURE MODES**14.1 Automating undefined metrics**

A model can extract a number perfectly and still produce a wrong report when the label, boundary or unit is undefined. Complete the metric contract before automation.

14.2 Treating AI output as evidence

A generated calculation explanation or narrative is a presentation layer. Preserve the source records, formula and review evidence independently.

14.3 Silent method drift

Emission factors, taxonomies, portfolio mappings and models change. Version them, approve effective dates and control restatements.

14.4 Missing represented as zero

Zero is a measurement. Missing is a data state. The system should never substitute zero for missing without an approved estimation method and explicit label.

14.5 False precision

Estimated activity, geospatial classifications and forecasts carry uncertainty. Displaying more decimal places does not improve evidence. Report precision consistent with source and method.

14.6 Boundary mismatch

Asset, fund, financial-control, operational-control and value-chain boundaries can differ. Store each boundary and test comparability before aggregation.

14.7 Double counting

Ownership structures, financed-emissions attribution, renewable instruments, carbon credits and impact claims can overlap. Use identity and rights records to prevent repeated inclusion.

14.8 Outcome promotion

Outputs can be mistaken for outcomes, and outcomes for contribution. Enforce the claims ladder and require evidence at each transition.

14.9 Narrative imbalance

Generative drafting can select positive facts and omit adverse effects, uncertainty or unresolved exceptions. Balanced-claim tests and human review should cover both positive and negative evidence.

14.10 Model bias and coverage gaps

Document types, languages, geographies, sensors and stakeholder groups may perform differently. Report subgroup performance and require abstention outside validated coverage.

14.11 Weak segregation of duties

One operator should not alter mappings, clear exceptions and release disclosures without independent control. Privileges and approvals must follow the control design.

14.12 Assurance theatre

An internal badge or model validation does not equal independent assurance. Use precise labels for review and assurance status.

14.13 Public claims without citation

Every material external statement should link to an approved metric or claim record. Unsupported text should remain blocked.

14.14 Shadow AI

Analysts may use unapproved tools to summarise confidential reports. Approved workspaces, data-classification controls, model inventory and training are required.

15. NINETY-DAY ADOPTION ROADMAP**Days 0-15: mandate and inventory**

- appoint executive sponsor, metric owners, data owners and release authority;
- identify applicable frameworks, covenants and report destinations;
- inventory assets, entities, sources, models and existing controls;
- select two bounded high-volume use cases;
- define privacy, security, assurance and retention constraints;
- establish value-ledger zeroes for unsupported financial fields.

Days 16-30: metric and evidence design

- build the priority metric catalogue and claims ladder;
- create identity rules for entity, asset, meter, investment and document;
- define source hierarchy, quality vector and exception taxonomy;
- design evidence packets and deterministic calculations;
- select representative test data, including adverse and low-quality cases;
- agree proposed acceptance thresholds and decision rights.

Days 31-45: build and shadow

- implement immutable source capture and controlled normalisation;

- configure extraction, mapping and anomaly services;
- preserve model version, prompt or configuration and citations;
- run shadow processing beside the current workflow;
- label every error and measure reviewer effort;
- prohibit external release from the shadow system.

Days 46-60: controlled pilot

- release one internal use case to a bounded population;
- require review of material fields and all low-confidence output;
- measure accepted packets, rework, exception burden and retrieval time;
- conduct privacy, access, incident and change-control tests;
- perform an assurance-readiness walkthrough;
- stop or narrow use when thresholds fail.

Days 61-75: harden and map reporting

- resolve root causes in identity, definition and source quality;
- approve factor, method and taxonomy versions;
- implement framework-specific disclosure mappings;
- add claims-register and narrative-support tests;
- document operating procedures, fallbacks and decommissioning;
- re-run the representative evaluation after every material change.

Days 76-90: bounded release

- approve the production boundary and named operators;
- retain human approval for judgement and external reporting;
- execute a complete period-close or reporting-cycle rehearsal;
- reconcile approved metrics to the disclosure output;
- update the value ledger with observed evidence;
- schedule drift, access, control and assurance reviews.

Ninety-day ESG data and AI adoption roadmap




Figure 8. Ninety-day ESG data and AI adoption roadmap

Source: Matchpoint implementation framework; governance context [14,15,22-25].

16. GOVERNANCE AND CONTROL CHECKLIST

Domain	Minimum control	Evidence
accountability	named sponsor, metric owner, data owner, reviewer and release authority	approved RACI and delegation
criteria	standard, regulation, covenant and method versions identified	criteria register
metric	definition, boundary, period, unit and claim class controlled	metric contract
source	population, identity, access and retention controlled	source register and reconciliation
model	intended use, owner, version, evaluation and monitoring recorded	model card and evaluation report
calculation	deterministic formula and factors reproducible	calculation specification and rerun
exception	severity, owner, due date and resolution retained	exception log
narrative	claims register, citation test and balanced presentation	approved claim packet
change	mapping, factor, rule and model updates approved	change ticket and test evidence
security	least privilege, logging, encryption and incident response	access review and security test
privacy	purpose, consent, minimisation and sensitive-data controls	privacy review
assurance	scope and status labelled precisely; evidence accessible	assurance pack and conclusion
value	baseline, accepted output, costs and unsupported zeroes preserved	value ledger

17. CLAIMS REGISTER AND LIMITATIONS

17.1 Supported propositions

The cited standards and official materials support the following propositions:

- investor-focused sustainability disclosures require governance, strategy, risk-management, and metrics and targets information [1,2];
- digital taxonomies enable structured tagging and machine-readable analysis [3,11];
- emissions and financed-emissions methods require defined boundaries, factors, attribution and data-quality disclosure [5-7];
- sustainability assurance requires suitable criteria, evidence and disciplined engagement performance [8,30];
- impact measurement should distinguish intent, outputs, outcomes, contribution and risk [17,18];
- digital MRV can integrate sensors, satellites, AI and registries, with interoperability and control requirements [19,20,27,28];
- AI risk management requires governance, measurement, testing, monitoring and documentation [22,23,25];
- UAE financial-sector sustainability frameworks emphasise systems, data governance, risk measurement and reporting [14,15].

17.2 Unsupported claims

This paper does not establish that an AI programme reduces a specific organisation's cash cost, increases revenue, prevents loss, improves investment returns or produces alpha. It does not establish assurance, regulatory compliance or impact additionality for any asset or portfolio. Those conclusions require approved observed evidence within the relevant boundary.

17.3 Empirical limitations

The paper is a standards-led operating framework. It is not a statistical study of AI performance, sustainability-reporting cost or investment outcomes. The two scenarios are unverified illustrative management assumptions. Proposed pilot thresholds require owner approval and empirical calibration.

17.4 Regulatory and professional limitations

Reporting and AI obligations vary by entity, jurisdiction, instrument and effective date. The EU and other frameworks continue to evolve. The paper is not legal, regulatory, accounting, audit, assurance, tax, environmental, engineering or investment advice. Qualified professionals should determine applicability and sign-off.

17.5 Technical limitations

Extraction, classification, entity resolution, anomaly detection, geospatial analysis, forecasting and generation can fail. Performance can change with source type, language, geography, model version and operating context. Human review, deterministic controls and fallback procedures remain necessary.

17.6 Data limitations

Infrastructure and private-asset data may be incomplete, delayed, estimated or confidential. Portfolio aggregation can conceal asset-level variation. External factors and taxonomies change. The evidence graph should retain all limitations rather than smoothing them away.

18. CONCLUSION

AI can improve ESG and impact measurement when it is placed inside a controlled evidence system. The system begins with decision lens, criteria, metric identity, boundary and source provenance. It uses deterministic calculation for reproducible metrics. It assigns AI bounded tasks with representative evaluation, abstention and monitoring. It gives reviewers an evidence packet, preserves decision rights and maps only approved facts into reporting taxonomies and narratives.

For A4 infrastructure investors, the priority is the connection between physical assets, digital observations, financial ownership and disclosure. For A2 family-office allocators, the priority is preservation of manager-reported evidence alongside transparent house normalisation. Both audiences need explicit data-quality vectors, claims ladders and assurance status.

The commercial discipline is equally strict. Measure accepted evidence packets and observed workflow effort. Preserve implementation and run costs. Attribute cash, revenue, loss reduction or alpha only after the causal bridge is approved and observed. Until then, the value ledger records USD 0 for those fields.

The result is an operating model in which AI supports faster and more consistent evidence work while accountability remains visible. That is the foundation for investment-grade sustainability information.

REFERENCES

- [1] IFRS Foundation. 2023. IFRS S1 General Requirements for Disclosure of Sustainability-related Financial Information. <https://www.ifrs.org/issued-standards/ifrs-sustainability-standards-navigator/ifrs-s1-general-requirements/>
- [2] IFRS Foundation. 2023. IFRS S2 Climate-related Disclosures. <https://www.ifrs.org/issued-standards/ifrs-sustainability-standards-navigator/ifrs-s2-climate-related-disclosures/>
- [3] IFRS Foundation. 2024. IFRS Sustainability Disclosure Taxonomy. <https://www.ifrs.org/issued-standards/ifrs-sustainability-taxonomy/>
- [4] IFRS Foundation. 2025-2026. IFRS Sustainability Disclosure Taxonomy Update - Amendments to Greenhouse Gas Emissions Disclosures. <https://www.ifrs.org/projects/work-plan/ifrs-sustainability-disclosure-taxonomy-update-amendments-ghg-emissions-disclosures/>
- [5] GHG Protocol. 2004. A Corporate Accounting and Reporting Standard, revised edition. <https://ghgprotocol.org/corporate-standard>
- [6] GHG Protocol. 2011. Corporate Value Chain (Scope 3) Accounting and Reporting Standard. <https://ghgprotocol.org/scope-3-calculation-guidance-2>
- [7] Partnership for Carbon Accounting Financials. 2025. Global GHG Accounting and Reporting Standard Part A: Financed Emissions, third edition. <https://carbonaccountingfinancials.com/standard>
- [8] International Auditing and Assurance Standards Board. 2024. International Standard on Sustainability Assurance 5000, General Requirements for Sustainability Assurance Engagements. <https://www.iaasb.org/consultations-projects/sustainability-assurance-issa-5000>
- [9] Global Reporting Initiative. 2021-2025. GRI Universal Standards and consolidated GRI Standards. <https://www.globalreporting.org/standards/>
- [10] Taskforce on Nature-related Financial Disclosures. 2023. Recommendations of the Taskforce on Nature-related Financial Disclosures, version 1.0. <https://tnfd.global/publication/recommendations-of-the-taskforce-on-nature-related-financial-disclosures/>
- [11] EFRAG. 2024. ESRS Set 1 XBRL Taxonomy and Explanatory Note. <https://www.efrag.org/en/projects/esrs-xbrl-taxonomy/concluded>
- [12] European Commission. 2026. Commission adopts revised sustainability reporting standards to reduce administrative burdens for EU businesses while maintaining high-quality disclosures. https://finance.ec.europa.eu/news/commission-adopts-revised-sustainability-reporting-standards-reduce-administrative-burdens-eu-2026-07-03_en
- [13] European Commission. 2026. Corporate sustainability reporting. https://finance.ec.europa.eu/financial-markets/company-reporting-and-auditing/company-reporting/corporate-sustainability-reporting_en
- [14] UAE Sustainable Finance Working Group. 2024. Principles for Sustainability-Related Disclosures for Reporting Entities. <https://rulebook.centralbank.ae/en/entiresection/46>
- [15] Central Bank of the UAE. Current through 2026. Climate-related Financial Risk Management Regulation. <https://rulebook.centralbank.ae/en/rulebook/climate-related-financial-risk-management-regulation>
- [16] International Capital Market Association. 2024. Handbook: Harmonised Framework for Impact Reporting for Green Bonds. <https://www.icmagroup.org/assets/documents/Sustainable-finance/2024-updates/Handbook-Harmonised-Framework-for-Impact-Reporting-June-2024.pdf>
- [17] International Finance Corporation. 2019-current. Operating Principles for Impact Management. <https://www.ifc.org/en/our-impact/impact-investing-at-ifc>
- [18] Global Impact Investing Network. Current. IRIS+. <https://iris.thegiin.org/>
- [19] World Bank and Carbon Markets Infrastructure Working Group. 2025. Technical Guidance Note on Standardizing Digital MRV in Carbon Markets: System Evaluation Criteria and Hotspots Assessment. <https://openknowledge.worldbank.org/entities/publication/397c4e52-445a-4cf4-89df-f2e61373a524>
- [20] United Nations Environment Programme, International Methane Emissions Observatory. Updated 2026. Methane Alert and Response System. <https://www.unep.org/topics/energy/methane/methane-alert-and-response-system-mars>
- [21] United Nations Framework Convention on Climate Change. Current. Transparency and Enhanced Transparency Framework resources. <https://unfccc.int/Transparency>
- [22] National Institute of Standards and Technology. 2023. Artificial Intelligence Risk Management Framework 1.0. <https://doi.org/10.6028/NIST.AI.100-1>
- [23] National Institute of Standards and Technology. 2024. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. <https://doi.org/10.6028/NIST.AI.600-1>
- [24] European Commission. Current through 27 July 2026. AI Act implementation and AI Omnibus timeline. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
- [25] Organisation for Economic Co-operation and Development. Updated 2024. OECD AI Principles. <https://oecd.ai/en/principles>
- [26] World Wide Web Consortium. 2013. PROV-DM: The PROV Data Model. <https://www.w3.org/TR/prov-dm/>
- [27] World Bank. 2022. Digital Monitoring, Reporting, and Verification Systems and Their Application in Future Carbon Markets. <https://openknowledge.worldbank.org/handle/10986/37622>
- [28] United Nations Environment Programme, International Methane Emissions Observatory. Current. Eye on Methane: MARS approach and methodology. <https://methanedata.unep.org/mars-approach-methodology>

[29] IFRS Foundation. Current. Supporting materials for IFRS Sustainability Disclosure Standards. <https://www.ifrs.org/supporting-implementation/supporting-materials-for-ifrs-sustainability-disclosure-standards/>

[30] International Auditing and Assurance Standards Board. 2025-2026. ISSA 5000 fact sheet, implementation guide and materiality FAQs. <https://www.iaasb.org/publications/issa-5000-fact-sheet>

APPENDIX A. MINIMUM ESG METRIC RECORD

Field	Example control response
Metric ID	stable code and version
Owner	accountable role and delegate
Criterion	standard or policy with version
Claim class	footprint, outcome, contribution, risk or narrative
Boundary	entity, asset, period, geography and value chain
Source population	expected records and reconciliation rule
Unit and factor	base unit, conversion, factor issuer and vintage
Transformation	extraction, mapping, aggregation and formula
Estimation	method, affected population, uncertainty and replacement plan
Model involvement	service, model version, evaluation and confidence
Validation	deterministic checks and exception status
Review	reviewer, date, evidence and conclusion
Assurance	precise scope and status
Disclosure	target report, taxonomy tag and narrative claim
Retention	source and evidence-pack period

APPENDIX B. PILOT SCORECARD

Measure	Definition	Result state
source coverage	source records received / required population	observed only
material-field precision	correct accepted extracted fields / accepted extracted fields	observed test
material-field recall	correct extracted material fields / labelled material fields	observed test
citation fidelity	supported accepted claims / accepted claims	observed test
entity-link error	incorrect material links / links reviewed	observed test
first-pass acceptance	packets accepted without material rework / packets reviewed	observed pilot
exception burden	exception hours / total workflow hours	observed pilot
reproducibility	metrics independently reproduced / metrics tested	observed test
unsupported-claim rate	unsupported material claims / material claims reviewed	observed test
incident count	security, privacy, disclosure or control incidents	observed pilot
cash cost reduction	approved avoided cash spend	USD 0 until evidenced
revenue	approved incremental contribution	USD 0 until evidenced
loss reduction	approved avoided loss	USD 0 until evidenced
alpha	approved net performance attribution	USD 0 until evidenced

APPENDIX C. ASSURANCE-READY EVIDENCE PACK

1. reporting criterion and version;
2. metric contract and boundary approval;
3. complete source-population reconciliation;
4. immutable source records and access log;
5. extraction output with evidence spans;
6. identity and mapping approvals;
7. deterministic formula, factors and rerun result;
8. estimates, uncertainty and replacement plan;
9. model version, representative evaluation and monitoring result;
10. exception log and resolution evidence;
11. reviewer sign-off and segregation-of-duties evidence;
12. taxonomy mapping and disclosure output;
13. restatement and change-control record;
14. claims-register entry for every material narrative statement;
15. precise internal-review or independent-assurance status.

APPENDIX D. GLOSSARY

A2. Matchpoint ICP for family-office CIOs and heads of alternatives.

A4. Matchpoint ICP for infrastructure and digital-infrastructure investors.

AI service. A bounded extraction, classification, entity-resolution, anomaly-detection, geospatial, forecasting or generative component.

Assurance. An engagement performed against suitable criteria under an applicable assurance standard; internal review is separately labelled.

Claim class. The controlled category of a statement, such as footprint, financed emissions, output, outcome, contribution, risk or narrative.

Data-quality vector. Separate attributes for provenance, coverage, completeness, freshness, method, control, uncertainty, assurance and comparability.

Digital MRV. Digitally enabled monitoring, reporting and verification that integrates controlled physical and digital evidence.

Evidence graph. Linked records for source, identity, transformation, calculation, review, approval and disclosure.

Financed emissions. Emissions attributed to financial activities under an approved method such as PCAF.

Impact. Change in an outcome experienced by people or planet, assessed with contribution and risk rather than output volume alone.

Materiality. A criterion-specific determination of information significance; financial materiality and impact materiality remain distinguishable.

Metric contract. The controlled definition, boundary, unit, method, source, owner and reporting mapping for one metric version.

Provenance. Information about the entities, activities and agents involved in producing or influencing data.

Taxonomy mapping. Approved relationship between a metric version and a reporting element or datapoint.

TEVV. Test, evaluation, verification and validation of an AI system.

Unverified illustrative management assumption. A worked input created only to demonstrate logic; it is not observed client or Matchpoint evidence.