

Net-Positive Management: A Governance, Risk, and Compliance Meta-Framework

Johan A. Pereira O.

johan@telos.org

[TELOS](#)

Abstract

Organizations implementing multiple management-system standards face structural fragmentation: independent scopes, duplicated evidence, and parallel audit cycles that erode strategic coherence and degrade governance into checklist compliance. This whitepaper introduces Net-Positive Management (NPM), a governance, risk, and compliance (GRC) meta-framework that integrates existing standards without replacing them. NPM is organized by four structural commitments: a teleological criterion derived from Polman and Winston's Net Positive standard; a logical gate of five ethical invariants grounded in Ellul's ethics of self-limitation and Jonas's ethics of responsibility; a Domain-Driven Design architecture (Evans, Brandolini) that models governance over stakeholder impact surfaces rather than organizational hierarchies; and a dynamic causal analytics pipeline (DCCA) operating under Pearl's structural causal models and Gama's concept-drift adaptation. The framework produces a traceable chain from strategic purpose through Bounded Contexts, Ethical Performance Indicators (EPIs), and Behavioral Event Signals (BES) to audit evidence. NPM materializes in seven artifacts constituting a Minimum Viable Compliance baseline and eight phases orchestrating their production and continuous verification. The whitepaper formalizes the ethical gate, the MCDA weight architecture, the DCCA four-control loop, and the operationalization of invariant standards, including extraction limits and the reasonable-third-party test. The framework is standards-agnostic and jurisdiction-adaptable.

Keywords: governance, risk, compliance, GRC, meta-framework, ethical invariants, self-limitation, causal analytics, Domain-Driven Design, MCDA, concept drift, ISO integration, Net Positive, Hans Jonas, Jacques Ellul, structural causal models

Key Concepts and Definitions

- **Efficiency** (Ellul, 2021). The optimization of means relative to ends involves evaluating outcomes based on the ratio of results to resources consumed. Ellul's critique is not that efficiency is undesirable, but that when efficiency becomes the supreme criterion of evaluation, the question of which ends are worth pursuing is displaced by the question of how to pursue whatever ends are already in place. In NPM, efficiency is treated as a subordinate value: legitimate

within the boundaries set by purpose and pathological when it replaces purpose as the governing criterion.

- **Purpose.** The reason an organization exists is expressed as the value it commits to creating for the stakeholders and systems it affects. NPM distinguishes between purpose as a criterion (a filter for what the organization should do) and purpose as a metric (a measure of how well it does it). The framework uses purpose exclusively as the former.
- **Technique** (Ellul, 2021). Technique is not synonymous with technology; it is the logic of optimization itself, applied to every domain from manufacturing to management to governance. When technique operates without constraint, it converts ends into means for further optimization. In Ellul's analysis, technique operates as an autonomous system that follows its own internal logic, displacing the question of purpose with the question of optimization.
- **Self-limitation & non-power** (Ellul, 2021; Jonas, 1984). The deliberate, architectural refusal to exercise power that is technically available but ethically impermissible. Ellul frames this concept as the "ethics of non-power," not powerlessness, but the conscious refusal to transform organizational power into manipulation, extraction, or capture of moral judgment. Jonas grounds the obligation, mentioning "responsibility is a correlate of power and must be commensurate with the latter's scope and that of its exercise" (1984, p. x). Self-limitation is constitutive of legitimate power. In NPM, self-limitation is a structural constraint encoded in the logical gate.
- **Net Positive** (Polman and Winston, 2021). The commitment that the organization will give back more than it takes to every stakeholder at every scale. In NPM, it is used as a teleological criterion (a filter for materiality decisions) rather than as a metric to be optimized.
- **Ethical Performance Indicator (EPI).** A lagging, outcome-based measure of whether an ethically material result has been achieved. EPIs are defined by the governance body through Phase F3 of the NPM pipeline and are anchored to the ethical gate. Full treatment is in the section "EPI and BES."
- **Behavioral Event Signal (BES).** A leading, activity-based measure of whether the upstream behaviors expected to produce an EPI outcome are occurring at the expected cadence and quality. Full treatment is in the section "EPI and BES."
- **Bounded Context** (Evans, 2004). The unit of domain modeling in Domain-Driven Design: a region of the organization with internally coherent language, rules, and data. Used in NPM to organize governance around stakeholder impact surfaces rather than departmental hierarchies.
- **Logical gate.** The architectural prerequisite that every candidate, control, policy, metric, or model must satisfy before entering the Compliance Core. Constituted by the five ethical invariants defined in the section "The Ethical Invariants."
- **Compliance Core.** The gated set of obligations, risks, controls, and evidence that constitutes the demonstrable governance commitment of the organization.

Only candidates that pass the logical gate, the traceability test, and the proportionality test are admitted.

- **North Star.** A direction-setting criterion that filters materiality rather than being optimized as a metric. The distinction between criterion and metric is structurally important: a criterion tells you whether a candidate control, policy, or investment belongs in the Compliance Core; a metric tells you whether a candidate performs well once inside it.

Background

Taking the International Organization for Standardization as a reference, we can see how each management system standard is published and treated independently. ISO 37000:2021 establishes governance principles. ISO 31000:2018 defines risk management. ISO 37301:2021 specifies compliance management systems. ISO/IEC 27001:2022 covers information security management systems (ISMS). ISO/IEC 42001:2023 covers artificial intelligence management systems (AIMS). ISO/IEC 27701:2025 covers privacy information management systems (PIMS). Each standard is internally coherent. None is designed with knowledge of the others' operational internals.

The practical consequence for an organization implementing several of these standards is the emergence of parallel governance structures: a scope statement per standard, a risk register per standard, a Statement of Applicability (SoA) per standard, and an audit program per standard. Evidence is duplicated, scopes overlap, responsibilities collide, and the total documentation burden grows faster than the governance value it produces. The organization ends up optimizing each standard's local checklist while losing sight of the question that should sit upstream of every standard: *what is this organization for, and at what cost to the systems that sustain it?*

Tactical compliance standards tend to degrade, under operational pressure, into binary compliance/non-compliance assessments. A control is either in place, or it is not. A clause is either satisfied, or it is not. A record is either filed, or it is not. This binary framing is convenient for audit but pathological for governance, as it cannot distinguish between a control that protects something of substance and a control that produces evidence of its execution without protecting the organization.

The origins of this derive are not malicious, but rather the predictable response of busy teams to frameworks that measure ritual execution and are silent on purpose. Compliance activities end up following a ritualistic pattern that signals conformity while losing contact with the outcomes they were designed to produce.

The Causal Blindness Problem in Conventional Risk Governance

ISO 31000 defines risk as "the effect of uncertainty on objectives," a definition that admits both positive and negative effects. In practice, implementations of risk

© 2026 [TELOS](#) This whitepaper is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit creativecommons.org

management typically reduce to a probability-by-impact matrix in which each risk is assigned a coordinate, and risks in the upper-right quadrant receive attention. This represents a reduction in the shape of correlational reasoning; that describes which risks co-occur with which impacts without articulating the causal pathway from one to the other.

The limitation is compounded by two distinct failure modes. First, probability-by-impact matrices are structurally unable to represent the second order and feedback effects that characterize most real organizational failures; they treat each risk as an independent coordinate rather than as a node in a causal network where interventions on one risk propagate to others. Second, and independently, Chernov, Ayoub, Sansavini, and Sornette (2023) demonstrate through systematic analysis of industrial disasters that even when risk information does exist within an organization, governance structures and cultural dynamics routinely prevent it from reaching the decision-makers who could authorize corrective action. The authors document a communicative and political failure where workers at affected sites knew conditions were dangerous, but organizational silence—driven by hierarchical pressure, fear of retaliation, and diffusion of responsibility—ensured that critical risk signals never reached the people empowered to act on them.

These represent two adjacent problems, where the first is an epistemological limitation of the quantification instrument itself and the second is a socio-organizational pathology that no improvement to the instrument can resolve in isolation. A governance framework that addresses only one will fail at the other. NPM's response is to commit to structural causal reasoning (Pearl, 2013) as the default framing for governance analytics, replacing correlational risk matrices with causal models that can represent feedback loops and propagation effects; and it requires, through invariants such like accountability for externalities and auditability of the ethical limit, that risk information flow be architecturally guaranteed rather than left to managerial discretion.

These are elements that contribute to governance, as the human system by which an organization is directed, supervised, and held accountable for its purpose (ISO 37000:2021). Governance must be purpose-oriented, not merely oversight-oriented, and it should produce sustainable value. NPM builds on the top of that definition of governance, providing methodological guidelines for linking the purpose declaration to the day-to-day operations of the organization. The purpose must thus be visible in the decision-making under the governance authority and not just relied upon as a documentation statement.

Technological Drift

Ellul, in *The Technological Society* (2021), argues that technique—defined as the totality of methods rationally arrived at and having absolute efficiency in every field of human activity—operates as an autonomous system that follows its own internal logic. When efficiency becomes the supreme criterion, the substantive question of purpose is displaced by the procedural question of optimization. In other words, when means become ends; then ends themselves are converted into means for further optimization. The organization becomes, in Ellul's terms, a *technical system* that cannot ask itself why it exists because such is not a technical question.

A related but distinct dynamic compounds this problem: organizations under measurement pressure tend to optimize for the indicators themselves rather than for the outcomes the indicators were designed to represent. This is a well-documented pathology in the measurement literature (cf. Goodhart's Law; Campbell, 1979) that converges with Ellul's diagnosis. Where Ellul identifies the autonomy of technique as displacing purpose at the systemic level, measurement fixation displaces it at the operational level. We're in front of two mutually reinforcing patterns where a system that cannot question its own purpose will default to optimizing whatever metrics are already in place, producing governance without substance.

In terms of corporate GRC, the absence of an explicit ethic of self-limitation in a governance framework that is initially purpose-oriented will cause, over time and under operational pressure, a drift into metric-oriented optimization. NPM invariants and the logical gate at the heart of it are the architectural response to this drift by imposing constraints that technique cannot metabolize into efficiency targets. The constraints are defined in terms of ethical boundaries rather than performance thresholds.

In summary, the NPM framework proceeds from five observations about the contemporary GRC landscape:

1. The governance frameworks can be internally coherent and externally uncoordinated.
2. Tactical standards degrade into binary checklist governance under operational pressure.
3. Probability-by-impact risk frameworks are causally blind and cannot detect drift.
4. Governance standards exhibit a purpose vacuum where they ask for purpose but cannot enforce its traceability to operations.
5. Without explicit self-limitation, frameworks drift toward technocratic optimization of the measurable.

NPM is designed as a response to these observations. The rest of this document specifies what that response consists of.

Theoretical Foundations

The first building block of NPM is Polman and Winston's (2021) reframing of corporate purpose as Net Positive, identified by the commitment of the organization to give back more than it takes to every stakeholder at every scale (employees, suppliers, communities, customers, future generations, and the planet itself). The authors describe this as a North Star, or a direction-setting criterion that filters what counts as material.

Polman and Winston's five operating principles translate the North Star into decision guidance:

1. Accept responsibility for all impacts and consequences, intended or unintended.
2. Operate for long-term benefits to the company and society.
3. Generate positive returns for all stakeholders.
4. Treat shareholder value as an outcome, not a goal.
5. Collaborate to drive systemic change.

The distinction between North Star as criterion and North Star as metric is structurally important to NPM. A criterion tells you whether a candidate control, policy, or investment *belongs* in the compliance core. A metric tells you whether a candidate performs well once inside it. NPM uses Net Positive only as the former. A candidate that does not improve the economic position of the business only when simultaneously reducing or correcting negative effects on the wider system fails the criterion and is rejected at the gate. It's designed as a conjunctive criterion that separates NPM from both shareholder-value reasoning (which ignores the wider system) and pure ESG rhetoric (which can rationalize economic underperformance). Stakeholder and economical sustainability are both material for an organization where NPM is applied.

The second building block is an ethics of self-limitation drawn from Ellul (2021). Ellul argues that a technological society, left unchecked, converts all human activities into problems of efficiency and thereby loses the ability to reason about ends. The response is not to reject technology but to impose an architectural limit translated in the ethics of no-power: not powerlessness, but the explicit refusal to transform organizational power into manipulation, extraction, or capture of moral judgment. The framework does not depend on the virtue of any individual decision-maker: it encodes the limit in the data model and the pipeline structure, so that even a decision-maker who wanted to bypass the limit could not do so without leaving audit evidence of the bypass. NPM's commitment is that ethical limits *must* be architectural and cultural, and must be gated at the point where candidates enter the Compliance Core, not checked later as an afterthought.

The third building block extends Ellul's diagnosis into an explicit ethics of responsibility. Hans Jonas (1984) argues that the scale and irreversibility of modern technological power demand a new ethical framework, one in which responsibility is commensurate with power and directed toward the future. Jonas formulates a new

categorical imperative: "Act so that the effects of your action are compatible with the permanence of genuine human life" (1984, p. 11). In NPM perspective, Ellul diagnoses the problem (technique as an autonomous system) and Jonas provides the normative response: an ethics in which the powerful become custodians of every end-in-itself that falls under their rule (1984, p. 130).

Jonas's concept of responsibility is differentiated from the contractual reciprocity that underlies most corporate governance. From Jonas point of view, responsibility is a non-reciprocal relation: the powerful are responsible for the vulnerable precisely because the vulnerable cannot reciprocate. This asymmetry (responsibility as a correlate of power, not of contract) provides the philosophical grounding for NPM's requirement that organizations account for externalities inflicted on parties who have no seat at the governance table.

NPM also considers the duty to know as defined by Jonas: "recognition of ignorance becomes the obverse of the duty to know and thus part of the ethics that must govern the evermore necessary self-policing of our outsized might" (1984, p. 8). An organization that does not know the consequences of its actions is not innocent but negligent. This concept drives the auditability requirement: ethical limits must produce evidence not because audit is convenient, but because ignorance of consequences is itself an ethical failure when the power to investigate exists.

In summary, NPM synthesizes these sources as follows: Polman and Winston supply the teleological criterion (what the organization is for), Ellul supplies the diagnostic (what will happen without constraint), and Jonas supplies the normative architecture (why constraint is obligatory and what form it must take). The five ethical invariants that follow are the operational encoding of this synthesis.

The Ethical Invariants

NPM encodes the ethics of no-power as five non-negotiable invariants. Each invariant is framed as a minimum, not a moral ceiling. An organization with more demanding ethics can add invariants, but the framework forbids removing or weakening the five.

1. **Purpose Primacy.** Every automation, control, metric, or model must demonstrate traceability to the organizational purpose and to the Net Positive North Star. A candidate that can be traced only to operational efficiency fails this principle. The operational test is whether the candidate has a documented chain from Bounded Context through strategic objective, to purpose declaration.
2. **Prohibition of Manipulation and Judgment Capture.** Decision methods that rely on hidden persuasion — dark patterns, internal propaganda, undisclosed nudging, or the substitution of decision support for decision capture — are forbidden. The invariant applies equally to anything produced by the organization, such as user-facing interfaces, internal management reporting, and automated decision systems. The operational test is whether a reasonable third

party, examining the mechanism, would recognize it as informing judgment rather than replacing it. (The operationalization of the "reasonable third party" standard is addressed in the section "Operationalizing Limits.")

3. **Extraction Limits.** For every material sustainability topic, the organization must identify the relevant ecological and social limits — planetary, legal, sectoral — and set internal boundaries that keep its impact within them. Where a limit is at risk of being exceeded, the organization must maintain repair and regeneration plans. The operational test is whether, for each material topic, there exists a documented limit, a monitoring mechanism, and a contingency plan. (The operationalization of the "reasonable third party" standard is addressed in the section "Operationalizing Limits.")
4. **Accountability for Externalities.** Ethical evaluation must consider unintended and second-order effects, including those associated with externalized segments of the business (suppliers, contractors, downstream users). The operational test is whether the risk register includes externalized effects with owners, and whether at least one evaluation cycle per year addresses externalities explicitly.

NPM adopts the concept of responsibility from Jonas (1984) where "responsibility is a correlate of power and must be commensurate with the latter's scope and that of its exercise" (1984, p. x). Moreover, "the first and most general condition of responsibility is causal power, that is, that acting makes an impact on the world" (p. 90). Accountability for externalities is therefore an obligation: the organization must act on consequences it discovers, not merely disclose them.

5. **Auditability of the Ethical Limit.** The ethical limit must produce evidence. Decisions, exceptions, compensations, and overrides must be recorded in a form that allows continuous evaluation with verifiable results. The operational test is whether every ethical decision — including the decision to grant an exception to an invariant — leaves a persistent record that an auditor can retrieve without asking the decision-maker.

Taken together, the five invariants constitute the logical gate. A candidate control, policy, metric, or model is admissible to the Compliance Core if and only if it passes all five invariants simultaneously.

Purposeful and Stakeholder-oriented architecture

The ethical gate determines what the organization must protect and pursue. The mission, vision, and goals of an organization determine why. NPM requires organizations to design their businesses considering how those commitments are maintained across every domain of organizational activity, under operational pressure, over time.

This responds to the very real risk of architectural drift. An organization may begin with a clear sense of purpose, but if its governance structures are modeled around internal authority lines rather than around the parties it exists to serve, the architecture will, incrementally and without announcement, reorganize itself around what is administratively convenient rather than what is ethically required. The commitment here is to an architecture designed around the impact surface of the organization, meaning the domains in which customers, clients, users, or other affected parties experience the consequences of organizational decisions.

NPM adopts Domain-Driven Design as the reference discipline for this architecture, drawing on Evans (2004), who introduced Bounded Contexts as the basic unit of organization for complex domains, and Brandolini (2020), who extended the same reasoning to business architecture through the Business Grid. This responds to an instrumental adoption selected because of the possibility to enforce three properties that a purposeful architecture requires:

1. Governance follows impact, not authority. Bounded Contexts are organized around the value delivered to impacted parties in each moment, not around departmental boundaries. A single domain may span multiple departments and business elements (products, features, customer-facing relationships, etc.), or a single element may respond to multiple domains. Modeling over the organization chart loses the structure of stakeholder impact entirely, because the organization chart represents reporting lines, not consequence lines.
2. The operational map is discovered from practice, not deduced from abstraction. Event Storming produces the event catalog by running a workshop in which the people who perform the work describe the events that occur, the commands that trigger them, and the policies that constrain them. The result is a map of what the organization does to and for its stakeholders, not a reference-model projection of what it should do. This fidelity to actual operations is what keeps the architecture answerable to the ethical gate rather than to a compliance template.
3. A Single Source of Truth (SSOT). Once drawn, the business grid supports the whole organization's practices, including its GRC controls, mapped risks and obligations simultaneously. Without this convergence, each domain constructs its own partial picture of the organization, and none of those partial pictures represents the full surface of stakeholder impact. Fragmented scoping is one of the main mechanisms supporting how purpose-oriented governance degrades into standard-by-standard box-ticking.

EPI and BES: Lagging Outcomes and Leading Behavioral Signals

Kaplan and Norton (1992, 1996) established the architectural principle that any performance measurement system requires both lagging indicators, which confirm

whether a desired outcome has occurred, and leading indicators, which track the upstream activities expected to produce that outcome. The balanced scorecard operationalized this pairing across financial, customer, process, and learning perspectives, but it did not address a prior question: what determines whether an outcome is worth measuring in the first place¹. NPM introduces two constructs that extend the leading/lagging architecture into governance by anchoring it to the ethical gate:

- An Ethical Performance Indicator (EPI) is a lagging measure of an outcome that the ethical gate has identified as mattering in itself: emissions reduced, disputes resolved, externalities compensated, audit findings closed, etc. EPIs tell whether the thing the organization committed to protecting or pursuing has been protected or pursued. The ethical mark indicates it as traceable to a gate decision, which means the organization can explain why this outcome was selected for measurement and not merely that it was.
- A Behavioral Event Signal (BES) is a leading measure of an activity that is causally upstream of an EPI: the frequency of compliance training completions, the rate at which incident tickets are filed within policy, the velocity at which risk reviews are conducted, etc. BES are procedural indicators that tell whether the activities expected to produce the outcome are occurring at the expected cadence and quality.

The distinction exists to prevent a specific failure mode. When an organization substitutes leading indicators for laggards (translated to EPI/BES in governance), it has migrated from governance to checklist compliance: the training was completed, the review was conducted, the box was ticked, but the outcome the activity was supposed to produce is never verified. The inverse failure, demanding EPI evidence without instrumenting the BES that would explain a miss, produces accountability without diagnosis.

NPM requires that every strategic objective be instrumented with at least one EPI and at least one BES and that the causal link between them be explicit. The link is the hypothesis that performing activity X at cadence Y will produce outcome Z. When the BES signals are healthy, but the EPI does not move, the hypothesis is wrong and the activity design must be revisited. When the EPI moves but the BES signals are absent or degraded, the outcome is occurring for reasons the organization does not understand, which is a different kind of governance risk.

¹ Their later work on strategy maps (Kaplan & Norton, 2004) partially addresses causal linkage by articulating the hypothesized relationships between perspectives, but the strategy map takes the strategic objectives as given. It organizes the logic of how objectives relate to each other while it does not impose a filter on which objectives qualify for inclusion or on what grounds.

© 2026 [TELOS](#) This whitepaper is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit creativecommons.org

EPIs are defined by the governance body during Phase F3 (KPI/KRI catalog) of the NPM pipeline. The process requires:

- **Strategic objective identification (Phase F0):** the governance body declares what the organization exists to achieve and protect.
- **Materiality determination:** for each objective, the governance body identifies the ethically material outcomes, those that the ethical gate requires the organization to track.
- **Indicator design:** for each material outcome, the governance body (supported by the compliance function and domain experts) defines the EPI: its formula, measurement cadence, data source, owner, and evidence requirements.
- **Gate validation:** each proposed EPI is tested against the five invariants. An EPI that cannot be traced to purpose (I1), that could be gamed through manipulation (I2), or that cannot produce audit evidence (I5) is rejected.

The governing body is accountable for the selection of EPIs. The compliance function is responsible for their design and instrumentation. Domain owners are responsible for data provision. Internal audit verifies that the EPIs remain anchored to the ethical gate over time. This allocation is reflected in the RACI matrix.

The Compliance Core

The Compliance Core is defined by three tests that every candidate must pass:

1. **Ethical acceptability test.** Does the candidate satisfy all five invariants?
2. **Traceability test.** Can the candidate be traced to at least one strategic objective and one Bounded Context?
3. **Proportionality test.** Is the candidate's operational cost proportionate to the risk it addresses and the value it protects?

A candidate that fails any one of the three is not a candidate. It is the architectural encoding of NPM's commitment that ethical limits cannot be optimized away by aggregating across other considerations.

Inside the gate, MCDA operates as the prioritization engine: among candidates that have passed the gate, which should be funded first, scaled first, and monitored most intensively? Outside the gate, no candidate is considered.

Operationalizing Limits

Two challenges recur in the application of the ethical invariants: the operationalization of the "reasonable third party" standard in Invariant I2, and the identification of social extraction limits in Invariant I3.

The "Reasonable" Standard

NPM requires that a mechanism pass the test of whether "a reasonable third party, examining the mechanism, would recognize it as informing judgment rather than replacing it." NPM addresses the operationalization through three mechanisms:

- **Legal precedent.** The "reasonable person" standard is the dominant standard in tort law, contract law, and professional regulation across common-law and civil-law jurisdictions. The "reasonable person" (common law), "bonus pater familias" (civil law), "ordenado empresario" (Spanish commercial law) is operationalized through centuries of case law, regulatory guidance, and professional norms. NPM inherits this operationalization rather than inventing its own.
- **Published reference norms.** The governance body is required to publish the reference norms against which "reasonable" is evaluated. These may include industry codes of conduct, regulatory guidance documents, professional standards, or the organization's own published values. The standard is therefore not "whatever a hypothetical person might think" but "what a person informed by these specific norms would conclude."
- **Audit trail.** The determination is auditable. When a mechanism is admitted to the Compliance Core, the assessment of its reasonableness is recorded, including the norms applied and the reasoning. When an auditor disagrees, the disagreement is recorded alongside the original assessment. The standard is operationalized by making the subjective judgment transparent, traceable, and contestable.

Social Extraction Limits

NPM under the Invariant I3 requires the identification of "relevant ecological and social limits" for every material sustainability topic. The ecological limits are comparatively straightforward: planetary boundaries (Rockstrom et al., 2009), regulatory thresholds, sectoral emission caps, but the social limits are harder.

NPM does not claim that social extraction limits can be identified with the same precision as ecological ones. What it requires is that the organization **engage with the question** rather than treat social impact as unmeasurable and therefore exempt from governance. The framework provides three operational mechanisms:

- **Stakeholder impact mapping.** The organization maps the parties affected by each Bounded Context and identifies, for each party, the dimensions of impact (labor conditions, data privacy, access equity, community displacement, attention extraction, mental health effects).
- **Reference frameworks.** For each dimension, the organization identifies the applicable reference standard: ILO conventions for labor, GDPR for data privacy, the UN Guiding Principles on Business and Human Rights for supply-chain

impacts, sectoral codes for industry-specific concerns, etc. The reference standard operationalizes the limit.

- **Thresholds of concern.** The limit is expressed as a threshold of concern rather than a bright-line boundary. For ecological limits, the threshold may be a quantitative cap (tons of CO₂, liters of water). For social limits, it may be a qualitative trigger: the point at which a reasonable assessment of stakeholder impact would require the organization to act. The governance body publishes the threshold and revises it in the annual strategic review.

The practical test for social extraction follows a five-step process:

1. Map the affected stakeholders and impact dimensions (users' attention, mental health, data sovereignty, content labor, community integrity);
2. Identify reference frameworks for each dimension;
3. Set thresholds of concern;
4. Instrument BES that monitor proximity to the thresholds;
5. Define contingency plans for threshold breach. The limits will be imprecise, contested, and evolving.

NPM's commitment is that they exist, are documented, are monitored, and are revised. An imprecise limit that is documented and monitored is seen as more governable than a precise outcome that is unmeasured because it was deemed unmeasurable.

Methodology Overview

NPM's analytical pipeline orchestrates the ingestion, validation, modeling, and scoring of data from all sources relevant to the Compliance Core, producing the composite risk and opportunity signals that feed the executive dashboard and the audit evidence repository.

The pipeline comprises a deliberately minimal set of stages, each of which is individually auditable and collectively sufficient to implement the DCCA loop. The names of the stages are generic labels; a concrete implementation may use different names.

Stage	Purpose	Auditable output
S0 Data Preparation	Source-specific pre-processing: format conversion, encoding normalization, deduplication, and structural alignment required before data can enter the ingestion layer. Scope depends on source heterogeneity.	Preparation log; source-to-staging mapping.
S1 Ingestion	Extract and load data from sources into the pipeline database.	Ingestion run records; source freshness timestamps.
S2 Correction	Apply deterministic corrections for known source quirks (type normalization, null handling, unit alignment).	Correction audit log.
S3 Validation	Check schema, completeness, consistency, and business-rule constraints.	Validation report; blocking-gate decision.
S4 Statistical Validation	Test for Simpson's paradox, distribution sanity, and outlier concentration before any modeling.	Statistical diagnostic report.
S5 Signal Modeling	Compute the individual criterion scores that will feed the MCDA composite. Includes any explainability outputs (SHAP, Partial Dependence Plots).	Criterion score table; explainability artifacts.
S6 Composite Scoring	Apply effective MCDA weights to produce per-entity composite scores.	Composite score table; per-entity attribution.
S7 DCCA Finalization	Run PSI drift detection; apply EMA weight update; persist weight snapshot and per-entity attribution.	DCCA snapshot; attribution rows.

The pipeline follows a decision flow, not a linear sequence. At each validation gate (S3 and S4), the pipeline evaluates whether the data meets published quality thresholds. If a gate fails, the pipeline halts, records the failure with diagnostic detail, and does not proceed to scoring. If both gates pass, stages S5 and S6 produce scores and attribution. At S7, a second decision branch occurs: PSI drift detection determines whether domain weights require forced recalibration ($PSI \geq 0.25$), are flagged for review ($0.10 \leq PSI < 0.25$), or proceed unchanged ($PSI < 0.10$). The EMA update runs regardless of the PSI outcome. Implementation documentation should render this logic as a decision-tree diagram to make the branching explicit; the table above describes the stages, not the control flow between them.




Figure 1. NPM Analytical Pipeline — Decision Flow with Blocking Gates and PSI Branching
During stages S0–S4 any failure blocks the pipeline and is recorded. Stages S5–S7 emit auditable outputs that become the inputs for executive decision-making and for the artifact catalog.

Multi-Criteria Decision Analysis (MCDA)

NPM uses MCDA (Ren, 2021) as the aggregation method inside each composite score. The choice is motivated by three MCDA properties:

- **Trade-offs are explicit:** A composite score is a weighted sum, and the weights are stated in policy. Auditors can inspect the weights; dissenters can challenge them; the governance council can adjust them.
- **Support of intra-domain ordering:** A governance council can set the relative importance of criteria within a domain without having to commit to a fixed domain-level influence.
- **Composability with dynamic weight adjustment:** The formulas below allow an adaptive loop to modify domain-level weights without breaking the policy-level ordering of criteria.

NPM organizes MCDA weights in three layers:

Layer	Notation	Source	Update cadence
Base criterion weights	$w_{base}(c)$	Governance council decision, published in policy	Static (policy review)
Domain weights	$w_{dom}(d)$	Observed data via DCCA	Every pipeline cycle (EMA)
Effective weights	$w_{eff}(c)$	Derived from the two layers above	Every pipeline cycle

The effective weight applied in the composite score is:

$$w_{eff}(c) = w_{dom}(d(c)) \times \frac{w_{base}(c)}{\sum_{k \in d(c)} w_{base}(k)}$$

Where $d(c)$ denotes the causal domain of criterion c . This formula has the following properties:

1. Intra-domain ordering is preserved. If $w_{base}(c_1) > w_{base}(c_2)$ for c_1, c_2 in the same domain, then $w_{eff}(c_1) > w_{eff}(c_2)$ regardless of the current value of w_{dom} .
2. Inter-domain emphasis is adaptive. The ratio $w_{eff}(c) / w_{eff}(c')$ for criteria in different domains scales with $w_{dom}(d) / w_{dom}(d')$.

When the governance has not defined differential intra-domain preferences, the default is equal base weights for all enabled criteria ($w_{base}(c) = \frac{1}{|d(c)|}$). A statistical validation gate may further restrict the enabled set per pipeline cycle, setting $w_{base}(c) = 0$ for criteria that fail to demonstrate discriminative power (measured as ROC AUC > 0.55, Mann-Whitney $p < 0.05$). The formula remains invariant under either regime.

As an example, suppose an organization's governance body defines three causal domains D1, D2, D3 with base weights:

Criterion	Domain	w_base
c1	D1	0.20
c2	D1	0.10
c3	D2	0.25
c4	D2	0.15
c5	D3	0.20
c6	D3	0.10

Domain base sums: $\Sigma D1 = 0.30$, $\Sigma D2 = 0.40$, $\Sigma D3 = 0.30$.

If the DCCA loop has adjusted domain weights to $w_{dom} = \{D1: 0.40, D2: 0.35, D3: 0.25\}$, the effective weights are:

$$\begin{aligned}
 w_{eff}(c_1) &= 0.40 \times \frac{0.20}{0.30} = 0.2667 \\
 w_{eff}(c_2) &= 0.40 \times \frac{0.10}{0.30} = 0.1333 \\
 w_{eff}(c_3) &= 0.35 \times \frac{0.25}{0.40} = 0.2188 \\
 w_{eff}(c_4) &= 0.35 \times \frac{0.15}{0.40} = 0.1313 \\
 w_{eff}(c_5) &= 0.25 \times \frac{0.20}{0.30} = 0.1667 \\
 w_{eff}(c_6) &= 0.25 \times \frac{0.10}{0.30} = 0.0833
 \end{aligned}$$

The intra-domain ordering is preserved across all three domains. The inter-domain emphasis reflects the current DCCA state.

Illustrative application. An organization operating a software-as-a-service (SaaS) business might define three causal domains for customer retention such as product usage, customer success engagement, and sales or commercial risk,

and place specific criteria in each domain. The method described above is the general one; any specific assignment of criteria to domains is an instance of the method, not the method itself.

The Dynamic Correlation-Causation Adjustment (DCCA) loop is NPM's operational answer to the question "how should weights be maintained when the underlying data-generating process is drifting?" DCCA is a discipline of four controls that, taken together, constitute the adaptive loop.

Control 1 Continuous Stability Monitoring (PSI)

Every pipeline cycle computes the Population Stability Index (PSI) between the current score distribution for each domain and a baseline distribution fixed at an earlier cycle. PSI is a symmetric divergence closely related to the Kullback-Leibler divergence:

$$PSI = \sum_{k=1}^K \left(p_c(k) - p_b(k) \right) \times \ln \frac{p_c(k)}{p_b(k)}$$

where $p_b(k)$ and $p_c(k)$ are the fractions of observations in bucket k for the baseline and current distributions, respectively, and K is the number of quantile buckets (typically 10).

PSI values are interpreted against three conventional bands:

Band	Interpretation	Action
$PSI < 0.10$	Statistically stable distributions.	No action.
$0.10 \leq PSI < 0.25$	Moderate drift. Distributions are diverging but not yet materially.	Log, flag for review, no forced recalibration.
$PSI \geq 0.25$	Significant drift. Distributions differ materially.	Force recalibration on this cycle; annotate the snapshot.

Two methodological commitments:

- **Baseline boundaries are fixed.** The quantile bucket boundaries come from the baseline distribution, not the current one. If boundaries are recomputed each cycle, PSI trivially approaches zero and drift detection is defeated.
- **Bands are published in advance.** PSI thresholds are set before any observation. It is part of a falsifiability commitment so that statements like "something has changed" can be proven wrong, rather than relying on a post-hoc judgment.

Control 2 Defined-Cadence Weight Recalibration (EMA)

Every pipeline cycle also runs an Exponential Moving Average update on the domain weights, regardless of whether PSI triggered drift. The update rule is:

$$w_{dom}^{new}(d) = (1 - \alpha) \times w_{dom}^{old}(d) + \alpha \times observed(d)$$

After all domains are updated, the vector is renormalized so that the sum of w_{dom}^{new} equals 1

The learning rate α is the single most consequential parameter of the loop. The reference implementation uses $\alpha = 0.10$, which gives:

- A half-life per observation of approximately $\frac{\ln 0.5}{\ln(1-\alpha)} \approx 6.58$ cycles. After seven cycles, half the weight comes from data observed in the last seven.
- A cap on single-cycle shift of $\alpha \times \max(|observed - old|) \leq 0.10$, even for violent signal changes. Large shifts across multiple cycles leave a clear audit trail.
- Stability in the practical range of $0.05 \leq \alpha \leq 0.15$. Lower values make the loop unacceptably slow. Higher values make it unstable under noise.

The observed signal $observed[d]$ is the current-cycle input to the EMA. NPM defines two modes for constructing it:

- **Phase 1 (coverage-based).** Used when outcome labels are not yet available (cold start, rare outcomes, delayed outcomes). The observed signal is the coverage rate of domain d , meaning the fraction of entities where at least one criterion in the domain has non-sentinel data. The semantic claim is that a domain producing non-trivial signal on more entities deserves more influence on the composite. It is a surrogate for causal influence, not causal influence itself; the framework labels it explicitly as such.
- **Phase 2 (correlation-based).** Used once outcome labels have accumulated. The observed signal is the Spearman rank correlation between the domain's raw aggregate score and the outcome label. Spearman is preferred over Pearson here because domain scores are typically bounded and non-Gaussian, the score-to-outcome relationship is often monotonic but non-linear, and rank correlation is robust to outliers.

As a summary, Phase 1 is Phase 2's scaffolding: it produces defensible weights before outcome history exists, so that the pipeline can deliver useful scores from the first cycle.

Control 3 Statistical and Data-Quality Validation

Before any data enters the weight-update computation, it passes a validation gate (stages S3 and S4 of the pipeline) which checks:

© 2026 [TELOS](#) This whitepaper is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit creativecommons.org

- Schema validity: required fields are present and well-typed.
- Completeness: missing-data rates per source are within published thresholds.
- Consistency: cross-source constraints, including referential integrity, hold.
- Distribution sanity: descriptive statistics on scalar fields are within credible bounds.
- Simpson's paradox: Aggregate trends are compared to subgroup trends; significant reversals are flagged.

Failure modes at this gate are explicitly non-silent. Silent degradation corrupts the weight snapshot and, through the EMA, propagates the corruption across several subsequent cycles before it is detected. The outcome is that a bad extract fails loudly.

Control 4 Change Logging and Audit Evidence

Every weight update produces a persisted snapshot row. Every per-entity attribution is tied to the snapshot that governed its calculation. A new cycle produces a new snapshot; no existing row is ever modified².

² This is the operational form of Invariant I5 (Auditability of the Ethical Limit).




Figure 2. DCCA Four-Control Loop — Validation, Drift Detection, Recalibration, and Audit

The Neutral Sentinel

A single scalar value (the reference implementation uses 50.0 on a 0–100 scale) is reserved as the neutral sentinel, or the default assigned to any criterion for which no data is available.

The sentinel solves a pervasive anti-pattern in composite scoring. The naive alternatives are:

- Assign a zero when data is missing. This punishes the absence of data as if it were bad news.
- Drop the criterion silently. This rebalances the composite behind the scenes without producing audit evidence of the adjustment.

Both distort the score. The sentinel alternative treats missingness as a first-class value in

the system:

- **Statistically neutral.** A criterion at sentinel contributes neither positive nor negative signal to the composite. The weighted-sum structure is preserved.
- **Operationally detectable.** Downstream code can distinguish “criterion is at sentinel” from “criterion is at the midpoint because the signal genuinely reads as neutral.”
 - A small tolerance band ($\epsilon = 1.0$ in the reference implementation) absorbs floating-point noise.
- **Required for Phase 1 coverage.** The coverage-rate signal is well-defined only when sentinel is explicit. Without it, “does this entity have data for this domain?” is a question we can’t compute.

The principal risk of sentinel-based design is misreading. A dashboard that shows “50 = neutral” without explaining that no data was collected will mislead every reader who interprets the number as a score. The remedy is labeling: any user-facing representation of a sentinel score must be tagged explicitly (“insufficient data”). The underlying scalar is an operational convenience; the label prevents it from being misread as meaning.

Audit Evidence Schema

NPM’s commitment to auditability is encoded in two persistent tables. Both are append-only; both are sized for auditability.

Weight Snapshots Table. With one row per pipeline cycle in which the weights were updated:

Column	Type	Purpose
snapshot_id	serial primary key (autoincremental, unique)	Immutable surrogate identifier referenced by attribution rows.
snapshot_date	timestamp	When the weights took effect.
w_[domain_k]	decimal	One column per causal domain; must sum to 1.0 across the row.
drift_detected	boolean	Whether PSI triggered drift on this cycle.
trigger_reason	varchar	scheduled, drift_detected, initial, or an operator-override label.
sample_size	integer	Number of entities in the batch used to compute the observed signal.
obs_[domain_k]	decimal	Raw observed signal per domain: coverage in Phase 1, correlation in Phase 2.
notes	text	Free-form audit narrative: PSI values, override justification, operator name.
created_by	varchar	Identity of the process or operator that wrote the row.
created_at	timestamp with	Row insertion timestamp, set by database trigger (DEFAULT

Column	Type	Purpose
	time zone	NOW()).
updated_at	timestamp with time zone	Last modification timestamp, set by database trigger on UPDATE. For append-only tables this column should never differ from created_at; its presence serves as a structural safeguard against silent mutation.

Some considerations about this table:

1. Snapshots are append-only. A new cycle produces a new row. No existing row is updated.
2. The created_at and updated_at columns provide an independent audit trail at the database level: if updated_at ever differs from created_at, a row has been mutated in violation of the append-only invariant, and the discrepancy is detectable.
3. The observed signal is stored, not just the resulting weights. If the weights are ever questioned, the signal that drove them is on record.
4. Free-form notes are a first-class field. Structured columns cannot anticipate every audit concern; the notes field absorbs the remainder.

Per-Entity Attribution Table. One row per entity per calculation date:

Column	Type	Purpose
attribution_id	serial primary key (autoincremental, unique)	Immutable surrogate row identifier. Exists to provide a stable, system-generated key independent of business semantics.
entity_id	varchar, not null	The business identifier of the entity scored: an account ID, asset ID, process ID, person ID, or whatever constitutes the unit of analysis in the organization's domain model. This is a business key that comes from the source system.
calculation_date	Date, not null	The date of scoring.
domain_score_[k]	decimal	Raw aggregated score per causal domain k.
domain_contribution_[k]	decimal	Fraction of the final composite contributed by domain k.
effective_w_[k]	decimal	Effective domain weight actually applied to this entity on this date.
snapshot_id	foreign key, not null	References the weight snapshot whose values were used.
created_at	timestamp with time zone	Row insertion timestamp, set by database trigger (DEFAULT NOW()).
updated_at	timestamp with time zone	Last modification timestamp, set by database trigger on UPDATE. For append-only tables this column should never differ from created_at; its presence serves as a structural safeguard against silent mutation.

Unique constraint: (entity_id, calculation_date). Exactly one attribution per entity per day. The distinction between attribution_id (surrogate) and entity_id (business key) is such that:

- attribution_id provides referential stability for joins and foreign keys.
- entity_id preserves the business-meaningful identifier that auditors and domain experts recognize.

Causation Structure Analyses

Phase 1 of the DCCA loop operates on correlational signals: identified from coverage rates, co-movement between criterion scores and outcomes, and EMA-smoothed weight adjustments. The question Phase 1 answers are "which domains co-move with the outcome?".

During Phase 2, the domain analysis is enhanced with three capabilities included in the pipelines: a falsifiable causal structure, an identification strategy for estimating causal effects from observational data, and a method for evaluating whether specific interventions produced their intended effects. These capabilities are derived from artifacts the organization has already produced in Phases F1 through F5, without requiring a separate causal modeling exercise. If Phase 1 shows correlation, Phase 2 explores causation.

The Business Grid, the Event Storming inventory, and the F3 metric registry already contain causal claims implicitly. Every BES-to-EPI link in the metric registry responds to a hypothesis: "performing activity X at cadence Y will produce outcome Z." Every inter-context dependency in the Business Grid is a directed structural claim: "Context A feeds Context B, not the reverse." Phase 2 conforms the formalization of these claims into testable structures. It is a progression of three methods, each appropriate at a different stage of organizational maturity. The three are complementary, so that each answers a different causal question, and each draws its inputs from artifacts the organization has already produced:

- **Method 1: Structural Equation Modeling (SEM).** Available immediately upon completion of F3. The SEM specification is derived from the Business Grid and the metric registry. It answers the structural question: "Is our hypothesized causal architecture consistent with the observed data?"
- **Method 2: Difference-in-Differences with Propensity Score Matching (DiD + PSM).** Available once the first F5 controls are deployed. DiD evaluates the effect of a specific intervention by comparing treated and untreated entities before and after the intervention. It answers the audit question: "Did this control produce its intended effect?"
- **Method 3: Bayesian Network with Interventional Semantics.** Available once the pipeline has accumulated sufficient cycles (defined as minimum 30 with stable PSI). The Bayesian Network promotes the SEM specification to a probabilistic model with full posterior distributions. It answers the decision question: "What is the probability of outcome Y if we intervene on variable X, given current uncertainty?"




Figure 3. Causation Structure Analyses: Maturity Progression³

³ High-level maturity progression of the Causation Structure Analyses. The Business Grid and F1--F3 artifacts seed Method 1 (SEM). As the organization deploys controls (F5) and accumulates pipeline cycles, Methods 2 and 3 become available independently. All three methods feed audit artifacts that close the governance loop through F0 strategic review.

Method 1: SEM Specification from the Business Grid

Input: A system of equations specifying which variables cause which, and the hypothesized direction. SEM tests whether a hypothesized structure is consistent with the observed covariance.

Derivation from existing artifacts

The F3 metric registry contains BES-to-EPI pairs with explicit causal hypotheses, where each pair is a term in SEM equations. The Business Grid's inter-context dependencies constrain permissible domain-level paths: a path from Domain A to Domain B is admissible when the Business Grid shows a dependency in that direction.

Extraction procedure:

1. For each EPI in the F3 registry, list every BES that is linked to it by the governance body. Each link becomes one term in the SEM equation for that EPI.
2. For each inter-context dependency in the Business Grid, note the direction. This constrains the domain-level path structure: no SEM path may contradict the Business Grid topology.
3. Assign initial coefficient estimates from the w_{base} weights in the MCDA table. These are the governance body's prior beliefs about relative importance.
4. Identify moderating variables from the F2 policy constraints. Any policy that constrains how a BES operates (e.g., "engagement is capped at five interactions per quarter per account") becomes a ceiling constraint or interaction term in the equation.

Output: Causal Specification Table

Column	Type	Purpose
spec_id	serial primary key	Immutable identifier for the causal specification row.
from_variable	varchar, not null	The BES or domain acting as the hypothesized cause.
to_variable	varchar, not null	The EPI or domain acting as the hypothesized effect.
direction	varchar, not null	positive, negative, or unknown. The hypothesized sign of the causal effect.
initial_coefficient	decimal	Initial estimate from  . Will be replaced by the SEM-estimated coefficient after fitting.
source_artifact	varchar, not null	F3 registry, Business Grid interface, or governance override. Traceability to the artifact that produced the hypothesis.

constraint	text	Policy constraint, ceiling, or interaction term that modifies the relationship. Null if unconstrained.
status	varchar	hypothesized, confirmed, rejected, or revised. Updated after SEM fit testing.
fit_pvalue	decimal	p-value from the SEM chi-square fit test. Null until the SEM is estimated.
created_at	timestamp with time zone	Row insertion timestamp.

SEM produces goodness-of-fit indices (CFI, RMSEA, SRMR, chi-square test) that evaluate whether the hypothesized structure is consistent with the observed data. If the fit indices reject the model (e.g., CFI < 0.90 or RMSEA > 0.08), the governance body's hypothesized causal structure is demonstrably inconsistent with reality. Then, the specification table is updated: rejected paths are marked, and the governance body must revise its F3 registry or its Business Grid dependencies accordingly. This revision feeds back into the next F0 cycle, thus closing the governance loop.




Figure 3.a. Method 1: SEM Pipeline⁴

Method 2: DiD + PSM

Input: A treatment group (entities that received the intervention), a control group (entities that did not), pre-intervention and post-intervention score observations, and a

⁴ Directed hypotheses are extracted from the F3 metric registry and constrained by the Business Grid topology. Initial coefficients come from governance-set w_{base} weights. The SEM fit test either confirms or rejects the hypothesized causal structure; rejected paths trigger an F0 strategic review.

credible parallel trends assumption.

Derivation from existing artifacts:

- Every control deployed in Phase F5 is an intervention.
- Event Storming (F2) identifies the moment the intervention occurs as a command-event pair.
- Treatment date: Event timestamp.
- Treatment group: Entities in scope.
- Control Group: Entities not in scope.
- Pre/Post Intervention Observations: Pipeline's persisted score snapshots (weight_snapshots and per-entity attribution tables).

Extraction procedure:

1. For each control in the F5 control library, identify the Bounded Context it operates in and the entities it affects.
2. From the per-entity attribution table, extract pre-intervention and post-intervention composite scores for both the treated entities and the untreated entities.
3. Test the parallel trends assumption: verify that the pre-intervention score trajectories for both groups moved together before the intervention. If they diverged, apply propensity score matching to construct a synthetic control group from the untreated pool, matching on observable characteristics (domain scores, entity size, tenure, or other covariates available in the attribution table).
4. Estimate the treatment effect: $(\text{post_treated} - \text{pre_treated}) - (\text{post_control} - \text{pre_control})$.

Output: Intervention Evaluation Record

Column	Type	Purpose
evaluation_id	serial primary key	Immutable identifier.
control_id	varchar, not null	References the F5 control library entry.
bounded_context	varchar, not null	The Bounded Context in which the control operates.
treatment_date	date, not null	The date the control was activated.
treatment_count	integer	Number of entities in the treatment group.
control_count	integer	Number of entities in the control group.

pre_period_cycles	integer	Number of pipeline cycles of pre-intervention data.
post_period_cycles	integer	Number of pipeline cycles of post-intervention data.
parallel_trends	boolean	Whether the pre-intervention trajectories satisfy the parallel trends assumption.
method	varchar	did, psm_did, or insufficient_data.
treatment_effect	decimal	Estimated causal effect of the control on the composite score.
confidence_interval_lower	decimal	Lower bound of the 95% confidence interval.
confidence_interval_upper	decimal	Upper bound of the 95% confidence interval.
p_value	decimal	Statistical significance of the treatment effect.
conclusion	varchar	effective, ineffective, or inconclusive.
created_at	timestamp with time zone	Row insertion timestamp.

Note: The most common obstacle is the absence of a control group. If an organization deploys all F5 controls simultaneously to all entities, there is no untreated group and DiD cannot be applied. The recommendation is staggered rollout: deploy each control to a subset of entities first, preserve a holdout group for at least four pipeline cycles, then evaluate. This is the same logic as a phased product rollout, applied to governance controls.




Figure 3.b.Method 2: DiD + PSM Pipeline⁵

⁵ Each F5 control deployment constitutes a natural experiment. Treatment and control groups are identified from the control's entity scope; pre- and post-intervention scores come from persisted attribution snapshots. If parallel trends fail, propensity score matching constructs a synthetic control. Results are persisted in the Intervention Evaluation Record for F7 audit.

Method 3: Bayesian Network from Accumulated Phase 1 Data

Input: Either an expert-specified directed graph with conditional probability tables, or sufficient multivariate time-series data for algorithmic structure learning.

NPM establishes algorithmic primacy over subjective heuristics. Expert-derived data is categorized as Reference Data and shall serve exclusively as an initial baseline or an audit benchmark for model validation. It shouldn't be considered an authoritative source once algorithmic processing is active and producing sufficient time-series data.

Derivation from existing artifacts

Every pipeline cycle produces scores per entity, composite scores, domain weights, PSI values, and per-entity attribution rows. After N cycles, this constitutes a multivariate time-series dataset where the joint distribution across criterion is observable.

Extraction procedure:

1. Use the SEM specification as the prior graph structure.
2. Feed the accumulated criterion-score time series into a structure learning algorithm (PC algorithm, GES, or score-based methods). The algorithm will confirm, reject, or discover edges the governance body have not yet hypothesized.
3. Where the learned structure agrees with the SEM specification, confidence is high. Where it disagrees, flag the discrepancy for the governance body: the data suggests a causal path that the Business Grid does not reflect, or vice versa.
4. Estimate conditional probability tables from the data. These replace SEM's point-estimate coefficients with full posterior distributions, enabling probabilistic interventional queries via do-calculus.
5. Apply do-calculus to answer governance questions: "If we force criterion X to increase by one standard deviation, what is the posterior distribution of the composite score change?"

Once the data supports it, the Bayesian Network constitutes the promotion of the SEM specification to a probabilistic model.

Transition criterion: minimum 30 pipeline cycles with stable PSI (below 0.10).

- This threshold is set to ensure the data-generating process is stationary enough for structure learning to produce reliable results.

Feedback to the Business Grid

When structure learning discovers edges that contradict the Business Grid, it constitutes a governance signal. The Business Grid represents the organization's

understanding of its own causal structure. If the data contradicts that understanding, the organization's understanding is assumed incomplete. The discovered discrepancy must be surfaced to the governance body and resolved in the next F0 review cycle: either the Business Grid is revised to reflect the discovered dependency, or the discovered dependency is explained by a confounding variable that the structure learning algorithm cannot observe. Both resolutions are documented in the causal specification table.




Figure 3.c.Method 3: Bayesian Network Pipeline⁶

In summary, the Business Grid is the seed for all three methods. The causal hypotheses are already embedded in the BES-to-EPI links and inter-context dependencies that

⁶ The SEM prior structure from Method 1 is promoted to a probabilistic model once 30 pipeline cycles with stable PSI have accumulated. Structure learning confirms or discovers causal edges; discrepancies with the Business Grid are surfaced to F0 for resolution. Conditional probability tables enable do-calculus interventional queries for governance decision support.

© 2026 [TELOS](#) This whitepaper is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit creativecommons.org

Phases F1 through F3 produce. Phase 2 asks the organization to formalize the causal claims it has already made and subject them to statistical scrutiny.

Explainability: SHAP, PDP, and the Pearl Commitment

NPM commits to explainable analytics as a precondition for admissibility in the Compliance Core. A model that cannot be explained cannot be audited; a model that cannot be audited cannot satisfy Invariant I5.

Two techniques are specified at the method level:

- **SHAP** (Lundberg and Lee, 2017). Shapley values provide a unified, additive attribution of a model's prediction to its input features. SHAP satisfies three desirable properties (local accuracy, missingness, and consistency) that make it the current default for tabular model explainability.
 - NPM recommends SHAP summaries at the model level (which features drive predictions overall) and SHAP waterfall plots at the individual level (which features drove a specific entity's score).
- **Partial Dependence Plots (PDP) and Individual Conditional Expectation (ICE)**. These visualize the marginal effect of a single feature on the model's prediction, holding others fixed.
 - SHAP is local and additive while PDP/ICE complements it in the global and functional levels.

These techniques answer the question "which features drove this score?" They do not answer the question "which features *cause* the outcome?" The distinction is essential. NPM treats SHAP and PDP as correlational explainability and requires that any claim of causality be supported independently, following Pearl's (2013) structural-causal-model framework.

Pearl's framework requires:

1. A directed acyclic graph (DAG) that encodes the causal hypotheses in structural form.
2. An identification strategy: back-door adjustment, front-door adjustment, or instrumental variables, that shows how to estimate a causal effect from observational data given the DAG.
3. A falsification test. An intervention or a natural experiment that is capable of refuting the causal hypothesis.

NPM does not automate Pearl's framework; what it does is mark the boundary: every claim in the executive dashboard is labeled either *correlational* (SHAP, PDP, DCCA Phase 1) or *causal* (DCCA Phase 2 with corroborating SCM evidence). A claim that cannot be supported at the level it is presented is downgraded or retracted.

Guarding Against Statistical Pitfalls


The methodology is designed to avoid five common pitfalls:

- **Correlation mistaken for causation:** the Phase 1 / Phase 2 distinction explicitly labels the coverage signal as a surrogate. The executive dashboard distinguishes correlational claims from causal claims.
- **Simpson's paradox:** stage S4 runs subgroup vs aggregate trend comparison. Reversals are flagged before composite scoring.
- **Multiple comparisons:** when multiple metrics are tested against an outcome, the pipeline applies Bonferroni or Benjamini–Hochberg correction. No metric is admitted on raw p-value alone.
- **Survivorship bias:** coverage rate uses the full active entity set in the denominator, not the set with complete data. Without this commitment, the rate trivially approaches 1.0 and loses its usefulness as a Phase 1 signal.
- **False precision:** every persisted score, weight, and PSI value carries a documented decimal precision. Weights are persisted to six decimal places, contributions to four, user-facing scores to two. Reporting more precision than the input data justifies is a form of misrepresentation⁷.

Core Framework: Architecture, Phases, and Artifacts

The NPM Architecture

NPM is organized as a four-layer architecture, moving from purpose to evidence:



⁷ Just because a number has six decimal places does not mean all six digits are meaningful. The system stores weights to six decimals and contributions to four for computational accuracy but displays only two decimal places to end users. Showing more precision than the underlying data supports would imply a false sense of exactness: for example, claiming a churn risk of 67.3842% when the input data can only reliably distinguish 67% from 68%.

- Layer 0 is the teleological criterion
- Layer 1 is the logical gate
- Layer 2 is the orchestration layer that replaces fragmented ISO implementations
- Layer 3 is the analytical engine that keeps Layer 2 honest under drift.

Nothing passes into Layer 2 that has not passed Layer 1. Nothing is measured in Layer 3 that is not anchored in Layer 2. Nothing in Layer 0 is a metric; it is a criterion.

Phases F0 through F7

The NPM pipeline is organized into eight phases. Each phase produces auditable deliverables and explicit decisions. The deliverables are not optional: an organization that implements NPM must produce all of them. This is the condition under which the framework is self-auditing.

Phase	Name	Technique	Primary deliverable
F0	Direction and scope	Definition of purpose, objectives, materiality	Scope statement; strategic objectives with measurable targets
F1	Business map	Business Grid / Context Mapping (Brandolini)	Domain map; Bounded Contexts; interfaces; dependencies
F2	Event discovery	Event Storming workshop	Event / command / policy inventory
F3	KPI/KRI catalog	Metric design, EPI/BES split	Metric registry: definition, formula, cadence, owner, evidence
F4	Correlational and causal analysis	DCCA pipeline	Criterion scores; composite scores; weight snapshots; attribution; explainability artifacts
F5	Intervention design	Control design workshop	Control library; test procedures; ownership matrix
F6	GRC translation	Artifact authoring	Policies, procedures, RACI matrices, registers, executive dashboard
F7	Audit and continuous improvement	Internal audit; management review	Audit reports; non-conformance register; improvement plan

- **Scoping and Design stages:** Phases F0 through F3. They are human-led workshops and governance decisions. They produce the inputs that the analytical pipeline (F4) will consume.
- **Translation stage:** Phases F5 and F6. The outputs of F4 are converted into controls, policies, and dashboards that the organization operates on.
- **Continuous verification stage:** Phase F7. The cycle closes, and non-conformances feed back into the next iteration of F0 through F3.

The pipeline is designed so that F4 runs on a published cadence (typically weekly or monthly), while F0 through F3 and F5 through F7 run on longer cycles (quarterly or

annually). The key invariant is that no metric reaches executive-level visibility without traversing all eight phases at least once. A metric introduced outside the phase sequence is an unauthorized metric; it must be retroactively anchored to an F0 objective, an F1 context, an F2 event, and an F3 registry entry before it is admissible.

The Minimum Viable Compliance Artifact Catalog (A0–A6)

The artifacts produced by NPM are a minimum viable compliance (MVC) catalog: a small set of documents that, together, constitute a defensible compliance baseline. Each artifact is mapped to a reference ISO standard, so that a single artifact serves multiple audit purposes.

Code	Artifact	Purpose	Primary references
A0	Scope statement (ISMS/PIMS/AIMS) + stakeholder map	Delimit the boundaries of governance; identify dependencies and stakeholders; prevent perimeter creep	ISO/IEC 27001; ISO 37301; ISO/IEC 42001; ISO/IEC 27701
A1	Compliance obligations register	Maintain a living list of identified obligations; enable compliance risk evaluation and regulatory-change management	ISO 37301 (obligations; documented information)
A2	Policies (security, compliance, privacy, AI)	Set intent, objectives, and continuous-improvement commitments; enable executive alignment	ISO/IEC 27001; ISO 37301; ISO/IEC 27701; ISO/IEC 42001
A3	Risk methodology and criteria (acceptance, evaluation, treatment)	Ensure repeatability, comparability, prioritization; document control decisions with justification	ISO 31000; ISO/IEC 27001 Cl. 6.1.2–6.1.3; ISO 37301 compliance risk assessment; ISO/IEC 42001 AI risk assessment
A4	Statements of Applicability (security + AI)	Evidence coverage of necessary controls; rationalize exclusions with traceability	ISO/IEC 27001 Annex A; ISO/IEC 42001 Annex A
A5	AI system impact assessment register	Identify, evaluate, and address impacts of AI systems; maintain responsible AI governance	ISO/IEC 42001 (AI system impact assessment)
A6	Internal audit program and management review	Convert conformity into a cycle of verification and improvement; reduce non-conformance recurrence	ISO/IEC 27001 Cl. 9.2–9.3; ISO 37301 Cl. 9.2–9.3; ISO/IEC 27701 Cl. 9.2–9.3; ISO/IEC 42001 Cl. 9.2–9.3

Each artifact is versioned, owned, classified, and retained according to a common

© 2026 [TELOS](#) This whitepaper is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit [creativecommons.org](https://creativecommons.org/licenses/by-nd/4.0/)

documentary control template. The same obligation, the same risk, the same control can appear in multiple artifacts without duplication, because the artifact layer is a view over a shared compliance repository rather than a set of standalone documents.

Minimum Metadata per Artifact

Every artifact carries at minimum⁸:

- Artifact code (A0–A6)
- Title
- Document ID
- Version
- Date
- Status (Draft / Under review / Approved / Obsolete)
- Owner
- Approver (top management)
- Classification (Public / Internal / Confidential / Restricted)
- Scope (ISMS / PIMS / AIMS / CMS)
- Repository URL
- Retention and disposition
- Change history (version, date, author, reason)

References Integration Map

The ISO standards referenced by NPM are integrated through the Compliance Core. The roles in the integrated system are:

Standard	Role in NPM
ISO 37000:2021	Governance frame: purpose, direction, accountability. Defines <i>what</i> governance is.
ISO 31000:2018	Risk frame: how uncertainty over objectives is evaluated and treated. Defines <i>how</i> risk is reasoned about.
ISO 37301:2021	Compliance management system: how obligations are identified, assessed, and evidenced.
ISO/IEC 27001:2022	Information security management system: how information assets are protected.
ISO/IEC 42001:2023	AI management system: how AI systems are governed, risk-assessed, and impact-assessed.
ISO/IEC 27701:2025	Privacy information management system: how personal information is governed (anchored to GDPR and equivalent regimes).

⁸ The metadata is the contract that guarantees accountability and it's by itself an auditable record.

Standard	Role in NPM
SOC 2 Type II	External assurance scheme for operational effectiveness of controls. Integrated as evidence in the same repository, not as a parallel program.
EU AI Act (Regulation (EU) 2024/1689)	Jurisdictional reference for AI system classification and risk obligations; anchored alongside ISO/IEC 42001 in Artifact A5.
GDPR (Regulation (EU) 2016/679)	Jurisdictional reference for personal-data processing; anchored alongside ISO/IEC 27701 in Artifacts A1 and A2.

The design commitment is that an organization compliant with NPM is compliant with each of the underlying standards as a byproduct. Certification becomes a natural consequence of a well-designed Compliance Core, rather than the primary object of compliance activity.

The RACI Model

NPM separates governance from management in line with ISO 37000. Governance defines purpose, values, and expectations, supervises the organization, and holds it accountable; management executes objectives within the parameters set by governance. Delegation does not transfer accountability.

The minimum RACI for the core NPM activities is:

Activity	Governing body / Top management	Compliance function	CISO / ISMS manager	DPO / PIMS owner	AI governance lead	Internal audit
Define purpose, values, stakeholder expectations	R/A	C	C	C	C	I
Approve policies (compliance, security, privacy, AI)	A	R	R	R	R	I
Register obligations and regulatory changes	A	R	C	C	C	I
Risk evaluation and treatment (security, AI)	A	C	R	C	R	I
Internal audit; reporting to management	A (receives results)	C	C	C	C	R

The governing body remains accountable (A) for every activity. Delegation to

operational functions is in the responsible (R) column. The internal audit function is the only role with independence from operational delivery, which is what makes its findings credible.

Implementation Guidance

For an organization with no existing standardized implementation, the recommended sequence is:

1. Run F0 (Direction and scope). Convene the governing body, declare the purpose, set initial strategic objectives with measurable targets, state the initial materiality scope. Produce A0.
2. Run F1 and F2 (Business map and Event discovery) in parallel. Host a Business Grid workshop with the executive team to map Bounded Contexts. Follow with Event Storming workshops in each context to produce the event inventory. The output is the foundation for every subsequent artifact.
3. Run F3 (KPI/KRI catalog). For each strategic objective, define at least one Ethical Performance Indicator (lagging, outcome-based) and at least one Behavioral Event Signal (leading, activity-based). Register them in A3 with formula, cadence, owner, and evidence source.
4. Stand up F4 (Analytical pipeline) in Phase 1 mode. Implement the pipeline stages, initialize DCCA with equal domain weights, and let the EMA loop adjust them from coverage signals⁹.
5. Draft A1 and A2 in parallel. The obligations register (A1) is informed by the jurisdictional analysis; the policy portfolio (A2) instantiates the minimum policy set for each management system.
6. Produce A3 and A4. The risk methodology is the governance document; the Statements of Applicability are the evidence that controls have been considered.
7. Run F5 (Intervention design). Convert F4 outputs into a control library with test procedures and ownership.
8. Run F6 (GRC translation). Publish the executive dashboard, the RACI matrix, and the operational runbooks.
9. Run F7 (Audit and continuous improvement). Plan the first internal audit against the newly published baseline. Findings feed back into F0 through F3 for the next cycle.

The full cycle is designed to be complete in a minimum of ten to twelve weeks in a small organization, and in parallel workstreams in a larger one. NPM has a commitment with traceability: every artifact must be anchored upstream (to purpose, context, event, and metric) before it is published.

⁹ There's no need to wait for outcome labels; Phase 1 produces sound scores from the first cycle.

Documentary Control Template

Every NPM artifact uses a common documentary-control template to ensure audit consistency. The template is independent of the artifact's substantive content.

Artifact code: A#
Title: [substantive title]
Document ID: [organizational document identifier]
Version: [semantic version]
Date: [issue date]
Status: Draft / Review / Approved / Obsolete
Owner: [named role or person]
Approver: [top-management role]
Classification: Public / Internal / Conf. / Restr.
Scope: ISMS / PIMS / AIMS / CMS
Repository / URL: [canonical location]
Retention: [retention period and disposition]
Change history:
Version | Date | Author | Reason

Cadences and Assurance

NPM operates on four cadences, each with its own deliverables and governance ritual:

- **Continuous.** Analytical pipeline (F4) runs on the published pipeline cadence, typically daily or weekly. Output: updated criterion scores, composite scores, DCCA snapshots, attribution rows.
- **Monthly.** Operational management review. Compliance function, risk owners, and CISO review the scorecard, the non-conformance register, and the in-flight remediation plans.
- **Quarterly.** Executive management review. The governing body reviews the scorecard, the effectiveness of controls, the status of objectives, and any proposed scope changes.
- **Annual.** Strategic review. Full re-evaluation of F0 (Direction and scope) considering the prior year's evidence. Initiates the next annual cycle of F0 through F7.

Internal audit operates on its own cadence, independent of the management cadences. The audit program (A6) defines frequency, methods, and reporting. SOC 2 Type II assurance, where applicable, is integrated into the same evidence repository rather than run as a parallel program.

A Worked Illustrative Example (Software-as-a-Service Provider)

To ground the method, consider a generic software-as-a-service (SaaS) provider. The example is illustrative, not prescriptive: every organization will instantiate NPM differently depending on its Bounded Contexts.

- **F0 Scope.** The provider declares its purpose as enabling customer sustainability

reporting; its strategic objectives include customer retention, platform reliability, and zero material data breaches. ISMS, PIMS, AIMS scopes are declared in A0.

- F1 Business map. Workshop produces Bounded Contexts including Product, Customer Success, Sales, Operations, and Legal. Interfaces and dependencies are mapped.
- F2 Events. Event Storming surfaces events such as “customer signed contract,” “customer logged first feature use,” “customer reported incident,” “account manager escalated risk,” “renewal negotiation started.”
- F3 Metrics. For the customer-retention objective, the provider defines one EPI (confirmed churn rate at 90 days) and three BES covering product usage, customer success engagement, and sales risk. Each metric has a formula, a cadence, an owner, and an evidence source.
- F4 Pipeline. The provider stands up the DCCA pipeline with three causal domains (Product, Customer Success, Sales), equal initial domain weights, and Phase 1 (coverage-based) observed signal. Sentinel value 50 is applied on a 0–100 scale for missing criteria. PSI bands are published at 0.10 / 0.25.
- F5 Controls. Controls include automated product-usage monitoring, customer-success health-score reviews, and sales escalation protocols. Each control is traced to the objective, the Bounded Context, and the metric it instruments.
- F6 Artifacts. The organization publishes A0 through A6. The executive dashboard presents the composite churn-risk score, the domain attribution, the PSI values, and the top SHAP features, labeled by correlational-versus-causal provenance.
- F7 Audit. The first internal audit confirms that every control in A4 is traceable to an obligation in A1 and a risk in A3. Non-conformances feed back into A6 and into the next F0 cycle.

This SaaS example is illustrative, not the subject of the whitepaper. The NPM framework applies equally to manufacturers, financial institutions, public-sector bodies, and any organization that faces the fragmentation problem described earlier.

Final considerations

- The five ethical invariants are minimums. An organization with a more demanding ethics can add invariants; NPM does not prevent an organization from adopting a weaker ethics if its jurisdiction permits. The gate is a floor that prevents obvious drift, not a ceiling that guarantees excellence. Organizations seeking genuine ethical leadership must build above the floor, not merely live on it.
- The causal promise of DCCA depends on the quality of the upstream data. Silent data degradation corrupts the weight snapshots and propagates across several cycles before PSI detects the drift. Control 3 of the DCCA loop (statistical and data-quality validation) is the operational safeguard, but it cannot detect every pathology.

© 2026 [TELOS](#) This whitepaper is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit creativecommons.org

- o Organizations implementing NPM must treat upstream data integrity as a first-class concern, with its own ownership, its own monitoring, and its own audit cycle.
- The domain-driven architecture assumes the organization either has or can produce a coherent business map.
 - o Organizations whose structure is radically shifting such as early-stage startups pivoting monthly, highly distributed collectives without stable boundaries, or organizations in active restructuring will struggle to apply NPM without first producing the map. In such cases, the recommended approach is to run a lightweight F1 workshop first, accept that the map will change, and revisit it on a short cadence until the organization stabilizes.
- Pearl's structural causal modeling requires either experimental data or strong assumptions about the underlying causal graph. Most governance contexts supply neither. NPM's response is layered: Phase 1 of DCCA is labeled as coverage-based, Phase 2 as correlation-based, and anything beyond that requires explicit SCM construction with stated assumptions and falsification tests. NPM does not solve the general problem of causal inference from observational data; it applies the best available methods within a governance context and labels the level of evidence honestly.
- The Reference integration map is anchored in European regulations (GDPR, the EU AI Act, DORA, and CSRD) because the framework was developed against that regulatory landscape. Organizations operating in other jurisdictions must adapt A1 (obligations register) and the jurisdictional references in A2 accordingly. The adaptation is mechanical in nature.

Conclusions

There is a comfortable lie at the heart of contemporary corporate governance: that compliance is a reliable indicator of integrity. Build enough controls, certify enough standards, produce enough documentary evidence, and the organization earns its right to operate. The Net-Positive Management (NPM) framework proposed in this whitepaper begins, precisely, by rejecting that lie.

The dominant GRC paradigm has fragmented governance into a portfolio of isolated cost centers, each optimizing its own checklist while the organization as a whole moves further from the question that actually matters: for whom does this enterprise exist, and at what cost to the systems that sustain it?

Using ISO as a reference, we've noticed:

- ISO 37000 offers a reference purpose-oriented answer at the governance layer but provides no method for translating the purpose into operations.

© 2026 [TELOS](#) This whitepaper is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit creativecommons.org

- ISO 37301 and ISO/IEC 27001 provide methods, but inside a binary compliance logic that systematically obscures causal complexity.
- ISO 31000 invites reasoning about uncertainty, but its implementations typically reduce risk to probability-and-impact matrices that are blind to the second-order effects that characterize real organizational failures.
- ISO/IEC 42001 and ISO/IEC 27701 broaden the perimeter into AI and privacy, but bring the same old vices proceduralism, audit-as-performance, the substitution of metric for meaning into the new territory.

NPM does not propose to replace any of these standards (or any other standard whatsoever). It proposes to constitute the connective tissue that is missing between them: a Compliance Core that treats governance as value infrastructure rather than as operational friction, governed by causal reasoning rather than by simple correlation, and bounded by an explicit ethical architecture (the logical gate) that prevents the framework itself from degenerating into the technocratic doctrine it was designed to overcome.

By integrating the Bounded Contexts of Domain-Driven Design with the causal modeling of Judea Pearl, the dynamic-weight adjustment of Gama-style concept-drift adaptation, and the multi-criteria prioritization of MCDA, NPM creates a traceable chain from governance decision to business outcome.

The taxonomy of artifacts (A0 through A6) and phases (F0 through F7) supplies the structural vocabulary that makes the chain documentable and auditable. The MVC architecture ensures that organizations can adopt the framework from a credible baseline, rather than facing the all-or-nothing paralysis that destroys most GRC transformation initiatives before they begin.

NPM is a GRC framework that encodes, as a non-negotiable architectural constraint, Ellul's warning about technological society: that when efficiency becomes the supreme value, an organization begins to measure only what is measurable and calls that what matters. The five ethical invariants (purpose primacy, prohibition of manipulation, extraction limits, accountability for externalities, and auditability of the ethical limit) are not recommended aspirations. They are architectural requirements. A control that does not pass the logical gate is not a control; it is a liability with documentation.

This whitepaper limits its contribution to the conceptual design and the formalization of the minimum viable compliance artifact catalog. As with any design proposal, it has limitations that deserve precision. The principal one is scope: NPM neither reproduces nor replaces the audit instruments, the technical controls, or the operational procedures that each reference standard defines in its official documents. That is not an omission; it is an architectural position. A framework that rewrote preexistent methods would lose exactly what distinguishes it: the ability to act as a neutral integration layer adaptable to different organizational contexts without being tied to a particular version of any standard.

© 2026 [TELOS](#) This whitepaper is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit creativecommons.org

Glossary

A0–A6 — Minimum Viable Compliance artifact codes for scope, obligations register, policies, risk methodology, Statement of Applicability, AI impact assessment, and audit program.

DCCA (Dynamic Correlation-Causation Adjustment) — NPM’s four-control methodology for maintaining dynamic weights under concept drift: PSI monitoring, EMA recalibration, statistical validation, audit logging.

AIMS (Artificial Intelligence Management System) — Management system as defined by ISO/IEC 42001:2023.

Bounded Context — The unit of domain modeling in Domain-Driven Design (Evans, 2004). A region of the organization with internally coherent language, rules, and data.

Business Grid — Brandolini’s visualization technique for mapping the full business across Bounded Contexts rather than the org chart.

BES (Behavioral Event Signal) — A leading, activity-based indicator of whether the upstream behaviors expected to produce an outcome are occurring.

Checklist governance — Degraded compliance in which evidence of process execution substitutes for evidence of purpose protection.

CMS (Compliance Management System) — Management system as defined by ISO 37301:2021.

Compliance Core — The architectural center of NPM: the gated set of obligations, risks, controls, and evidence that constitutes the demonstrable commitment of the organization.

EMA (Exponential Moving Average) — Recursive smoothing formula used by DCCA to recalibrate domain weights each cycle: $\text{new} = (1 - \alpha) \times \text{old} + \alpha \times \text{observed}$.

EPI (Ethical Performance Indicator) — A lagging, outcome-based indicator of whether an ethically material outcome has occurred.

Event Storming — Brandolini’s workshop method for discovering the events, commands, and policies that populate a Bounded Context.

F0–F7 — NPM pipeline phases: Direction and scope; Business map; Event discovery; KPI/KRI catalog; Correlational and causal analysis; Intervention design; GRC translation; Audit and continuous improvement.

GRC (Governance, Risk, and Compliance) — The integrated field comprising governance, risk management, and compliance management.

ISMS (Information Security Management System) — Management system as defined by ISO/IEC 27001:2022.

Logical gate — The architectural prerequisite that every candidate control, policy, metric, or model must satisfy before entering the Compliance Core. Constituted by the

© 2026 [TELOS](#) This whitepaper is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit creativecommons.org

five invariants.

MCDA (Multi-Criteria Decision Analysis) — The aggregation method used inside NPM’s composite scores. Supports explicit weight assignment, trade-off inspection, and intra-domain ordering.

MVC (Minimum Viable Compliance) — The minimal set of artifacts (A0–A6) that constitutes a defensible compliance baseline under NPM.

Net Positive — Polman and Winston’s (2021) teleological criterion: the commitment to give back more than is taken, to every stakeholder at every scale.

Neutral sentinel — The scalar value (e.g. 50 on a 0–100 scale) reserved as the default for criteria with no available data, preserving the additive structure of the composite score.

North Star — A direction-setting criterion that filters materiality rather than being optimized as a metric.

NPM (Net-Positive Management) — The governance meta-framework presented in this whitepaper.

PDP (Partial Dependence Plot) — Visualization of the marginal effect of a single feature on a model’s prediction, holding others fixed.

PIMS (Privacy Information Management System) — Management system as defined by ISO/IEC 27701:2025.

PSI (Population Stability Index) — Symmetric divergence measure used by DCCA to detect drift between a current score distribution and a baseline.

RACI — Responsible / Accountable / Consulted / Informed matrix for assigning roles to activities.

SCM (Structural Causal Model) — Pearl’s (2013) framework for representing causal relationships as directed acyclic graphs and estimating intervention effects.

SHAP (SHapley Additive exPlanations) — Lundberg and Lee (2017) method for attributing a model’s prediction to its input features using Shapley values.

SoA (Statement of Applicability) — The compliance artifact (Artifact A4) that evidences which controls have been selected and justifies exclusions.

VUCA — Volatile, Uncertain, Complex, Ambiguous. The operating context in which static governance assumptions fail and dynamic loops like DCCA become necessary.

References

AICAD Business School. (2025). *Unidad 3. El riesgo tiene que ver con el conocimiento.*

Brandolini, A. (2020). *Context Mapping on a Business Grid.* Avanscoperta Blog. <https://blog.avanscoperta.it/2020/04/21/context-mapping-on-a-business-grid/>

Brandolini, A. (2025). *Introducing EventStorming.* Leanpub. https://leanpub.com/introducing_eventstorming

Campbell, D. T. (1979). Assessing the impact of planned social change. *Evaluation and Program Planning*, 2(1), 67–90.

Chernov, D., Ayoub, A., Sansavini, G., & Sornette, D. (2023). *Averting Disaster Before It Strikes: How to Make Sure Your Subordinates Warn You While There Is Still Time to Act.* Springer Nature. <https://doi.org/10.1007/978-3-031-30772-0>

Directive (EU) 2022/2464 of the European Parliament and of the Council of 14 December 2022 amending Regulation (EU) No 537/2014, Directive 2004/109/EC, Directive 2006/43/EC and Directive 2013/34/EU, as regards corporate sustainability reporting (CSRD). *Official Journal of the European Union*, L 322. <http://data.europa.eu/eli/dir/2022/2464/oj>

Ellul, J. (2021). *The Technological Society.* Knopf Doubleday Publishing Group. (Original work published 1964.)

Evans, E. (2004). *Domain-Driven Design: Tackling Complexity in the Heart of Software.* Addison-Wesley.

Gama, J., Žliobaitė, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. *ACM Computing Surveys*, 46(4), 1–37. <https://doi.org/10.1145/2523813>

Goodhart, C. A. E. (1984). *Problems of Monetary Management: The U.K. Experience.* In *Monetary Theory and Practice.* Macmillan.

International Organization for Standardization. (2018). *ISO 31000:2018 — Risk management — Guidelines.* <https://www.iso.org/standard/65694.html>

© 2026 [TELOS](#) This whitepaper is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit creativecommons.org

International Organization for Standardization. (2021a). *ISO 37000:2021 — Governance of organizations — Guidance*. <https://www.iso.org/standard/65036.html>

International Organization for Standardization. (2021b). *ISO 37301:2021 — Compliance management systems — Requirements with guidance for use*. <https://www.iso.org/standard/75080.html>

International Organization for Standardization. (2022). *ISO/IEC 27001:2022 — Information security, cybersecurity and privacy protection — Information security management systems — Requirements*. <https://www.iso.org/standard/27001>

International Organization for Standardization. (2023). *ISO/IEC 42001:2023 — Information technology — Artificial intelligence — Management system*. <https://www.iso.org/standard/42001>

International Organization for Standardization. (2025). *ISO/IEC 27701:2025 — Security techniques — Extension to ISO/IEC 27001 and ISO/IEC 27002 for privacy information management — Requirements and guidelines*. <https://www.iso.org/standard/27701>

Jonas, H. (1984). *The imperative of responsibility: in search of an ethics for the technological age*. Univ. of Chicago Press

Kaplan, R. S., & Norton, D. P. (1992). *The balanced scorecard: Measures that drive performance*. Harvard Business Review, 70(1), 71–79. <https://hbr.org/1992/01/the-balanced-scorecard-measures-that-drive-performance-2>

Kaplan, R. S., & Norton, D. P. (1996). *The Balanced Scorecard: Translating Strategy into Action*. Harvard Business School Press.

Kaplan, R. S., & Norton, D. P. (2004). *Strategy maps: Converting intangible assets into tangible outcomes*. Harvard Business School Press.

López Cerezo, J. A. (2018). *La confianza en la sociedad del riesgo* (1ª ed.). Sello.

Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. *arXiv preprint*

© 2026 [TELOS](#) This whitepaper is licensed under the Creative Commons Attribution-NoDerivatives 4.0 International License. To view a copy of this license, visit creativecommons.org

arXiv:1705.07874.

<https://doi.org/10.48550/arXiv.1705.07874>

Patton, J. (2014). *User Story Mapping: Discover the Whole Story, Build the Right Product*. O'Reilly Media.

Pearl, J. (2013). *Causality: Models, Reasoning, and Inference* (2nd ed., reprinted with corrections). Cambridge University Press.

Polman, P., & Winston, A. S. (2021). *Net Positive: How Courageous Companies Thrive by Giving More Than They Take*. Harvard Business Review Press.

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation). *Official Journal of the European Union*, L 119. <http://data.europa.eu/eli/reg/2016/679/oj>

Regulation (EU) 2022/2554 of the European Parliament and of the Council of 14 December 2022 on digital operational resilience for the financial sector (DORA). *Official Journal of the European Union*, L 333. <http://data.europa.eu/eli/reg/2022/2554/oj>

Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act). *Official Journal of the European Union*, L 2024/1689. <http://data.europa.eu/eli/reg/2024/1689/oj>

Ren, J. (2021). *Multi-Criteria Decision Analysis for Risk Assessment and Management*. Springer International Publishing.

Verwijns, C., Schartau, J., & Overeem, B. (2021). *Zombie Scrum Survival Guide: A Journey to Recovery*. Addison-Wesley.

Wallerstein, I. M. (2011). *El moderno sistema mundial* (2nd expanded ed., P. López Máñez, Trans.). Siglo Veintiuno.