

Sustainable Sovereign Cloud Transformation: A Harmonised Framework Architecture for ESG-Compliant, Risk-Bounded Enterprise Cloud Operations in the European Union

Justice Opara-Martins, PhD
EU-EC Digital4Business (D4B) Advanced Tech Consortium
National College of Ireland (NCI),
Mayor Street Lower, IFSC, Dublin 1, Ireland

ABSTRACT

Enterprise cloud transformation in the European Union sits at the intersection of sustainability reporting (CSRD/ESRS), operational and cyber resilience regulation (NIS2, DORA) and AI governance (the EU AI Act). Existing maturity models treat these as separate compliance exercises rather than as one economically quantifiable transformation problem. This paper proposes a harmonised four-construct architecture: a Sustainable Sovereign Cloud Transformation Blueprint (SSCTB) for structural alignment, a Cloud Sustainability Maturity Model (CSMM) for diagnosis, a Sustainable Cloud Value Framework (SCVF) for economic quantification, and a Transformation Risk Index (TRI) as a bounded risk moderator. It formalises their relationships, and **proves that the proposed model is bounded and monotone**. A pilot Delphi panel and survey ($n = 20$) give directional support for the maturity–value relationship (H1), partial support for risk moderation (H2) and a first calibration of the cost-of-capital mapping (H3); the direct TRI effect is not significant, and the pilot is small, non-independent and developer-administered. **A multi-criterion evaluation maps each validity claim to its present evidence and to the study still required**. The contribution is a pilot-tested design architecture, not a fully validated model.

General Terms

Information Systems, Cloud Computing, Digital Transformation, Sustainable Technology, Enterprise Architecture, Cybersecurity and Operational Resilience, Regulatory and Technology Governance

Keywords

Cloud sustainability; enterprise cloud architecture; ESG compliance; maturity model; digital sovereignty; cost of capital; EU regulation; CSRD; NIS2; DORA engines

1. INTRODUCTION

European cloud-dependent organisations now manage a combined transformation problem rather than a set of independent compliance tasks. Corporate sustainability reporting under the Corporate Sustainability Reporting Directive (CSRD) and the European Sustainability Reporting Standards (ESRS) requires auditable environmental information [3,4]. The NIS2 Directive and the Digital Operational Resilience Act (DORA) extend cybersecurity, continuity, and third-party-risk obligations across critical digital operations [5,11]. The EU Artificial Intelligence Act adds governance requirements wherever cloud-based systems support regulated or high-risk AI use cases [12]. These

requirements converge on architecture decisions: workload location, data-residency controls, telemetry, vendor concentration, portability, and evidence production.

The resulting organisational problem is not merely that firms lack a cloud strategy. It is that sustainability, resilience, sovereignty, and economic justification are often owned by different functions and measured through non-comparable instruments. This fragmentation can produce duplicated assurance work, incomplete risk visibility, and an inability to express cloud-transformation priorities in a language shared by architecture, sustainability, security, and finance teams. Digital transformation research similarly identifies the need to connect technological change to organisational capabilities and outcomes rather than treating technology adoption as an isolated event [14].

1.1 Problem and Gap

EU enterprises face a compounding compliance burden. The Corporate Sustainability Reporting Directive and its reporting standards require granular, auditable environmental disclosure [1], [2]; NIS2 and DORA impose operational-resilience and third-party risk obligations on cloud-dependent organisations [3], [4]; and the AI Act introduces governance duties for AI systems increasingly embedded in cloud operations tooling [5]. Cloud architecture decisions (workload placement, vendor selection, telemetry governance) now sit inside this perimeter, yet sustainability, resilience and AI governance are typically pursued as separate initiatives with separate tooling, ownership and reporting lines. Digital-transformation scholarship, which frames transformation as strategic response to digital disruption [6], does not by itself supply an integrated, quantitative account of this regulatory convergence.

This fragmentation is argued to impose three costs: duplicated compliance effort, inconsistent risk visibility across domains, and an inability to express the investment case in terms a CFO or board can act on. The pilot evidence in Section 5.2.5 is consistent with the first two. **The gap is therefore integrative rather than additive**: what is missing is a single model that relates structural architecture, maturity, risk exposure and economic value within one internally consistent, testable formalism.

1.2 Aim, Objectives and Research Questions

The aim is to consolidate the structural, diagnostic, economic and risk constructs of EU enterprise cloud transformation into

one formally consistent and empirically testable architecture. Each objective answers one research question:

- **O1 / RQ1** (Section 3): consolidate the constructs into one architecture. *How can cloud sustainability maturity, transformation risk and economic value be formally and consistently related in a single quantitative model?*
- **O2 / RQ2** (Section 4): specify a fieldable mixed-methods protocol. *What is the minimum viable empirical protocol capable of testing that model's core proposition to a standard defensible in a high-impact journal?*
- **O3 / RQ3** (Sections 5–8): demonstrate, evaluate and draw implications. *What are the practical and policy implications for EU enterprises and member-state regulators?*

1.3 Contributions and Organisation

The paper makes three contributions. First, it specifies a dimensionally consistent value–risk formalism and **establishes its boundary and monotonicity properties analytically** (Section 3.3). Second, it provides a fieldable protocol and a pilot ($n = 20$), together with **an explicit evaluation of what the evidence does and does not establish** (Section 5). Third, it derives implications and a falsifiable research agenda (Sections 6–8). Section 2 sets out the theoretical and regulatory foundations, Section 3 the architecture, Section 4 the method, Section 5 the findings and evaluation, Section 6 the discussion, Section 7 the limitations and Section 8 the conclusions. The two data-collection instruments are provided as Supplementary Material S1 and S2.

2. THEORETICAL FOUNDATION AND RELATED WORK

Cloud sustainability and sovereignty are therefore treated as organisational capabilities, not only as infrastructure properties. A mature organisation should be better able to sense carbon and regulatory signals, seize optimisation and compliance opportunities, and reconfigure workloads, vendors, and controls.

2.1 Dynamic Capabilities and Cloud Transformation

The architecture draws on dynamic capabilities, an organisation's capacity to sense, seize and reconfigure resources under environmental change [7], to explain why cloud maturity should predict enterprise value under regulatory pressure. The resource-based view [8] complements this: dynamic capabilities explain how a capability is built and renewed, while the resource-based view explains why, once built, it is difficult to imitate and so defensible rather than merely present. **Applied here, regulatory convergence is the environmental shift to be sensed; integrated telemetry, sovereignty controls and exit plans are the reconfigured resources; and CSMM measures how far that reconfiguration has been institutionalised.**

A boundary condition applies. Dynamic capabilities theory presumes that resources can be reconfigured at a pace matching environmental change, whereas a five-level ladder is a discrete snapshot of a continuous, path-dependent process. CSMM is therefore a diagnostic cross-section, not a claim that organisations move through levels in fixed, uniform steps; this is why scoring is continuous (Section 3.2) and levels are used only for interpretability. **The logic also yields a falsifiable**

implication: where risk constrains reconfiguration (for example through lock-in), the value yield of an additional unit of maturity should fall. This is H2 (Section 3.4).

2.2 EU Regulatory Landscape

Table 1 maps each instrument onto the framework's constructs. The mapping is what allows the framework to operate as a single compliance-and-value model, and it is offered as a structural claim to be tested (Section 4), not as an established finding.

Table 1. Regulatory instruments and their framework mapping

Instrument	Core obligation	Primary framework mapping
CSRD [1] / ESRS [2]	Auditable sustainability disclosure	Regulatory Assurance (RA), Carbon Efficiency (CE)
NIS2 [3]	Cybersecurity risk management and incident reporting	TRI, particularly vendor-dependency and operational-disruption domains
DORA [4]	ICT and third-party risk in financial entities	TRI, particularly vendor-dependency and operational-disruption domains
AI Act [5]	Risk-based governance of AI systems	Operational Optimisation (OO), where AI-assisted tooling is in scope

2.3 Positioning Relative to Existing Standards and Maturity Models

CSMM is positioned relative to, not as a replacement for, existing management-system standards and maturity approaches. **ISO 14001 [9], ISO/IEC 27001 [10] and ISO/IEC 42001 [11] specify auditable management systems for environmental, information-security and AI governance respectively. They certify conformity of process; they do not grade an organisation on a continuous scale or relate that grade to economic value. They are therefore complements that supply evidence for RA, SR and OO scoring, not substitutes.** CSMM's distinguishing claim is the formalised linkage from maturity, through a bounded value score, to a cost-of-capital adjustment (Sections 3.3–3.4), a linkage typically asserted only qualitatively. This is a testable differentiator, not a superiority claim.

The contribution is not a claim that one maturity score replaces legal compliance. Rather, the architecture offers a common measurement interface through which an enterprise may identify where regulatory assurance, sustainability telemetry, sovereignty controls, and transformation risk intersect.


3. CONCEPTUAL FRAMEWORK ARCHITECTURE

Four constructs are specified and measured, and together form the architecture (Figure 1); a construct earns its place by being operationalised, not merely named.

- **SSCTB** (Sustainable Sovereign Cloud Transformation Blueprint): structural alignment of

- infrastructure, data-sovereignty perimeters and regulatory constraints.
- **CSMM** (Cloud Sustainability Maturity Model): a five-level ordinal diagnostic over continuous 0–100 scores (Section 3.2).
- **SCVF** (Sustainable Cloud Value Framework): a bounded composite economic scoring model (Section 3.3).
- **TRI** (Transformation Risk Index): a bounded risk-exposure moderator (Section 3.3).

Figure 1. Harmonised conceptual architecture: the structural (SSCTB), diagnostic (CSMM), risk (TRI) and economic (SCVF) layers each converge on measurable enterprise outcomes.



Construct roles and ordering. SSCTB is a precondition for meaningful scoring; CSMM diagnoses capability; TRI moderates how much of that capability converts into value; SCVF expresses the outcome and feeds the cost-of-capital mapping. **SCVF has two operational roles that must not be conflated.** SCV, as defined by Equation (1), is a *model-implied* value score computed from CSMM inputs and TRI. Perceived value (the four outcome items SCV1–SCV4 in Supplementary Material S2) is an *observed* outcome. Because Equation (1) contains the CSMM–TRI interaction by construction (Section 3.3), H1 and H2 must be tested against the observed outcome; regressing model-implied SCV on its own inputs would be partly tautological.

3.1 SSCTB: Structural Layer

SSCTB specifies four structural domains that should be satisfied before CSMM/SCVF scoring is meaningful: (i) jurisdictional data placement mapped against sovereignty requirements; (ii) a documented telemetry architecture feeding carbon, cost and risk metrics into a common data model; (iii) an explicit vendor-exit and portability plan; and (iv) governance ownership assigned jointly across sustainability, security and enterprise-architecture functions. SSCTB is a precondition checklist, not a maturity score.

3.2 CSMM: Diagnostic Layer

CSMM scores an organisation on five vectors (Carbon Efficiency, Regulatory Assurance, Sovereign Resilience, Operational Optimisation and Sustainable Innovation), each on a continuous 0–100 scale mapped to one of five ordinal levels (Table 2). Continuous scoring is what allows CSMM to combine arithmetically with TRI without a unit mismatch.

Level labels are nominal descriptors and imply no equivalence with generic capability-maturity ladders.

Table 2. CSMM maturity levels, characteristics and score bands

Level	Characteristics	Score band (0–100)
1 – Initial	Ad hoc, reactive; no unified telemetry; manual ESG reporting	0–20
2 – Defined	Documented policies; carbon/ESG metrics tracked but siloed	>20–40
3 – Optimised	Cross-domain telemetry integrated; automated dashboards; sovereignty controls emerging	>40–60
4 – Managed	Predictive, AI-assisted optimisation; standardised, interoperable architecture	>60–80
5 – Sovereign	Autonomous, self-healing, closed-loop operations; full jurisdictional control	>80–100

3.3 SCVF and TRI: Economic and Risk Layers

Earlier practitioner versions of the model computed $SCV = (CE + RA + SR + OO + SI) - TR$, combining five 1–5 variables

with a 0–100 risk index. **The composite then ranges from –95 to +25: the risk term numerically dominates and value can be negative.** The corrected model scores every input, including TRI, on a common 0–100 continuous scale and applies TRI as a bounded multiplicative discount:

$$SCV = [W_{CE} \cdot CE + W_{RA} \cdot RA + W_{SR} \cdot SR + W_{OO} \cdot OO + W_{SI} \cdot SI] \times (1 - TRI/100) \quad (1)$$

$$TRI = (RR + CR + VDR + ODR + SCR) / 5 \quad (2)$$

where the weights W are governance-configured and sum to 1.0, and RR , CR , VDR , ODR and SCR are the regulatory, cyber, vendor-dependency, operational-disruption and sovereignty-compliance risk domains, each scored 0–100. Writing M for the weighted maturity composite:

Proposition 1 (bounded value). If every input lies in $[0,100]$, weights are non-negative and sum to 1, then $M \in [0,100]$, $(1 - TRI/100) \in [0,1]$ and therefore $SCV \in [0,100]$. Moreover $\partial SCV / \partial M = 1 - TRI/100 \geq 0$ and $\partial SCV / \partial TRI = -M/100 \leq 0$: the marginal value of maturity falls linearly as risk rises. This is the structural counterpart of H2.

The earlier cost-of-capital adjustment, $\beta_{post} = \beta_{baseline} \times (1 - \ln(1 + SCV/5))$, is unbounded: the logarithmic term exceeds 1 once $SCV > 5(e - 1) \approx 8.59$, making β_{post} negative, an uninterpretable systemic-risk coefficient. It is replaced by a bounded exponential-compression form:

$$\beta_{post} = \beta_{baseline} \times [f + (1 - f) \cdot \exp(-k \cdot SCV/100)] \quad (3)$$

with beta-floor ratio $0 < f < 1$ and compression sensitivity $k > 0$.

Proposition 2 (bounded, monotone compression). For $f \in (0,1)$ and $k > 0$, the ratio $\beta_{post}/\beta_{baseline}$ is strictly decreasing in SCV , lies in $(f, 1]$ for every $SCV \geq 0$, and on $SCV \in [0,100]$ lies in $[f + (1 - f)e^{(-k)}, 1]$. Hence $\beta_{post} > 0$ always. **The floor f is an asymptote, approached only as $k \cdot SCV/100$ grows large; it is not reached at $SCV = 100$ unless k is large.** The resulting cost-of-equity relation retains its original form and now inherits a well-behaved β_{post} :

$$K_e = R_f + \beta_{post} \times (R_m - R_f) - \Delta_{ESG} \quad (4)$$

f and k are empirical parameters to be fitted against observed beta/ESG-score pairs (Section 4.4), not asserted from theory; Section 5.2.4 reports a first pilot calibration.

3.4 Hypotheses

- **H1:** CSMM maturity score is positively associated with observed value outcomes, controlling for sector and organisation size.
- **H2:** TRI negatively moderates the CSMM-to-value relationship: the marginal value of a maturity gain is smaller for high-risk organisations.
- **H3:** The value-to-cost-of-capital mapping is bounded and compressive: higher value is associated with equal or lower risk-adjusted cost of equity, never a negative systemic-risk coefficient. **Because Equation (3) is monotone by construction (Proposition 2), the empirical content of H3 lies in whether observed beta/ESG data are fitted by the compressive form with $k > 0$, not in monotonicity itself.**

4. METHODOLOGY

The primary theoretical contribution is architectural rather than causal. The paper shows how dynamic-capabilities reasoning

can be translated into measurable cloud practices while preserving distinctions among readiness, maturity, risk, and value [1,13]. The bounded equations also demonstrate that a practitioner-originated framework can be made dimensionally coherent and mathematically well-behaved without claiming that the equations themselves establish empirical validity.

4.1 Research Design

A sequential mixed-methods design is used: a qualitative Delphi phase (Supplementary Material S1) validates construct content and item wording, followed by a cross-sectional survey (S2) that tests H1–H3 through reliability analysis, confirmatory factor analysis and moderated regression. A longitudinal extension is specified in Section 8. **The illustrative scenarios of Section 5.1 are computational demonstrations of the model's arithmetic: hypothetical composites, not primary data. They are kept analytically separate from the pilot (Section 5.2) and from the planned full-scale study ($n \geq 100$).**

4.2 Delphi Panel Protocol

Following established Delphi practice [12], the panel comprises 15–20 experts spanning enterprise cloud architecture, sustainability/ESG reporting, EU regulatory policy and academic information-systems research. Panellists rate each candidate item for relevance, clarity and completeness on a 1–7 scale. An item is retained when the inter-quartile range (IQR) is ≤ 1.0 , at least 75% of the panel agree on relevance, and median relevance is ≥ 6.0 . The process stops after three rounds or earlier if ratings are stable across two consecutive rounds. Panel size, response rate per round and consensus statistics are reported, not only the final item list.

4.3 Quantitative Survey Instrument

The survey (S2) contains 34 items on a five-point agreement scale: five CSMM constructs of four items each (20), five TRI domains of two items each (10, protective items reverse-scored), and four perceived-value outcome items ($SCV1 - SCV4$). The target sample for the full-scale study is $n \geq 100$, stratified by sector and organisation size, drawn independently of the Delphi panel.

4.4 Analysis Plan

- **Reliability:** Cronbach's alpha [13] per construct, reported regardless of outcome (≥ 0.70 acceptable, ≥ 0.80 good, ≥ 0.90 excellent).
- **Dimensionality:** exploratory factor analysis on a random split-half and confirmatory factor analysis on the remainder, reporting standardised loadings and fit indices. **Two threshold tiers are used. A conventional tier (CFI, TLI ≥ 0.90 ; RMSEA, SRMR ≤ 0.08) screens the pilot. The stricter criteria of Hu and Bentler [14] (CFI, TLI ≈ 0.95 or above; RMSEA ≈ 0.06 or below; SRMR ≈ 0.08 or below) are the confirmatory standard for the full-scale study.**
- **Hypothesis tests (H1, H2):** moderated regression of the observed value outcome on CSMM, TRI and their interaction with sector and size controls, reporting exact coefficients, standard errors, p-values, 95% confidence intervals and effect sizes.
- **H3:** nonlinear regression of observed listed-company beta on SCV to estimate f and k in Equation (3).

- **Common-method bias:** Harman's single-factor test as a diagnostic of limited sensitivity, supplemented by procedural and statistical remedies.

5. FINDINGS AND EVALUATION

Section 5.1 demonstrates the proposed model computationally, Section 5.2 reports the pilot, and Section 5.3 evaluates what the combined evidence does and does not establish.

5.1 Model Behaviour: Illustrative Application and Sensitivity

Table 3 applies Equations (1) and (3) to three hypothetical composite profiles spanning the CSM range. The inputs are illustrative and measure no organisation. The final column uses the pilot-fitted f and k of Section 5.2.4 only to show the full chain from maturity to financing cost.

Table 3. Illustrative application of the corrected model (hypothetical profiles; model-implied values)

Illustrative profile	CSM M level (band)	Maturity input M	TRI	SCV	$\beta_{\text{post}}/\beta_{\text{baseline}}$
Early-stage adopter (hypothetical)	2 (>20–40)	32	58	13.4	0.95

Illustrative profile	CSM M level (band)	Maturity input M	TRI	SCV	$\beta_{\text{post}}/\beta_{\text{baseline}}$
cal mid-sized manufacturer)					
Mid-transformation (hypothetical public-sector estate)	3 (>40–60)	54	41	31.9	0.88
Advanced adopter (hypothetical financial-services firm)	5 (>80–100)	88	19	71.3	0.76




Figure 2. Model-implied behaviour of the corrected equations: (a) SCV against maturity at three TRI levels; (b) the beta ratio against SCV at the pilot-fitted and two illustrative values of k .

Model-implied curves computed from Equations (1) and (3) with pilot parameters $f = 0.052$, $k = 0.412$; $k = 1$ and $k = 2$ are illustrative sensitivities. Not empirical data.

Figure 2 makes two properties visible. In panel (a), because $\partial \text{SCV} / \partial M = 1 - \text{TRI} / 100$, a maturity point is worth 0.8 value points at TRI = 20 but only 0.2 at TRI = 80: the moderation hypothesis H2 is the empirical counterpart of this structural property. In panel (b), at the pilot-fitted $k = 0.412$ the beta

ratio falls only to 0.68 at SCV = 100, so the fitted floor $f = 0.052$ lies far outside the range the data can reach. k carries the empirical content of Equation (3); f is an extrapolative parameter, weakly identified until larger samples or observations at higher SCV constrain it.

5.2 Pilot Empirical Findings (n = 20)

5.2.1 Sample and Design Caveats

A first-round Delphi panel and a pilot administration of the survey were completed with the same cohort of 20 senior enterprise decision-makers (6 cloud/enterprise architects, 5 IT/infrastructure directors, 4 ESG/sustainability officers, 3 risk/compliance officers, 2 financial/investment directors; 85% with more than ten years' experience; Germany, France, the Netherlands, Ireland, Portugal and Italy). **Three caveats govern every result below.** (i) Panel and survey draw on the same individuals, so the two stages are not independent triangulation. (ii) $n = 20$ is underpowered for confirmatory factor analysis and moderated regression (Section 5.3), so all statistics are directional. (iii) The pilot was administered by the framework's developer. The independent full-scale study ($n \geq 100$) is designed to resolve all three; the results here are pilot evidence, not the pre-specified validation.

5.2.2 Measurement: Delphi Consensus and Reliability

Table 4 reports Delphi consensus against the thresholds of Section 4.2 and Cronbach's alpha against Section 4.4; Figure 3 shows the same results against those thresholds. Five of six construct groups were retained in Round 1. Sustainable Innovation (SI) fell short (IQR 1.25, 70% agreement, median 5.5), was revised toward cloud-provider-native, carbon-aware workload shifting on panellists' comments, and reached consensus in Round 2 (IQR 1.00, 80% agreement).

Table 4. Pilot measurement results by construct (n = 20)

Construct	Median relevance (1–7)	IQR	% agreement	Delphi status	Cronbach's α
Carbon Efficiency (CE)	6.5	1.00	85%	Retained, Round 1	0.84
Regulatory Assurance (RA)	7.0	0.50	95%	Retained, Round 1	0.91
Sovereign Resilience (SR)	6.0	1.00	80%	Retained, Round 1	0.88
Operational Optimisation (OO)	6.5	1.00	80%	Retained, Round 1	0.79
Sustainable Innovation (SI)	5.5	1.25	70%	Revised; consensus Round 2	0.76
Transformation Risk Index (TRI)	6.5	1.00	85%	Retained, Round 1	0.83
SCVF composite	–	–	–	–	0.87




Figure 3. Pilot measurement results: (a) Delphi agreement by construct against the 75% threshold; (b) Cronbach's alpha by construct against the 0.70 threshold.

Pilot Delphi consensus and Cronbach alpha as reported in the submitted manuscript (Tables 4 and 5), $n = 20$; thresholds from Sections 4.2 and 4.4.

All seven coefficients exceed 0.70 and RA exceeds 0.90. **With $n = 20$ each estimate carries wide sampling uncertainty, and an alpha above 0.90 on four items can also signal item redundancy; the coefficients are therefore directional.**

Panelists also rejected the equal-weighting default and recommended sector-stratified weights. For highly regulated sectors (financial services, healthcare), $RA \approx 0.30$, $SR \approx 0.30$ and $SI \approx 0.10$, with CE and OO sharing the remaining 0.30; for less-regulated, technology-native sectors, CE and OO ≈ 0.25 each. **This is a testable refinement of the governance-configured weights in Equation (1), to be validated on the full-scale sample before adoption as a default.**

5.2.3 Structural Fit

Table 5 reports the pilot fit indices against both threshold tiers of Section 4.4.

Table 5. Pilot confirmatory fit against two threshold tiers ($n = 20$)

Index	Pilot value	Conventional tier	Hu and Bentler tier [14]	Outcome
CFI	0.924	≥ 0.90	≈ 0.95	Conventional : met. Strict: not met
TLI	0.908	≥ 0.90	≈ 0.95	Conventional : met. Strict: not met
RMSEA	0.074	≤ 0.08	≈ 0.06	Conventional : met. Strict: not met
SRMR	0.068	≤ 0.08	≈ 0.08	Both met

The pilot indices pass the conventional tier but meet only SRMR under the stricter criteria. **They cannot be read as confirmatory for a more basic reason: the 30 CSMM and TRI items exceed the 19 degrees of freedom that 20 respondents provide for the sample covariance matrix (its rank is at most 19), so item-level maximum-likelihood CFA is not identified.** The analysis is therefore an internal-consistency screen, not structural evidence.

5.2.4 Hypothesis Tests

Table 6. Pilot hypothesis test results ($n = 20$)

Hypothesis parameter	Estimate	Significance	Status
H1: CSMM $\rightarrow$ value (unstandardised)	+0.642	$p < 0.001$	Supported (pilot)
H1: adjusted R^2	0.581	–	–
H2: direct CSMM effect	+0.512	$p < 0.005$	Significant
H2: direct TRI effect	–0.210	$p = 0.082$	Not significant

Hypothesis parameter	Estimate	Significance	Status
H2: CSMM $\times$ TRI interaction	–0.315	$p < 0.015$	Partially supported (pilot)
H3: baseline K_e (fixed assumption)	8.50%	fixed, not fitted	Input assumption
H3: beta-floor f (fitted)	5.20%	–	Fitted (pilot)
H3: compression sensitivity k (fitted)	0.412	$p < 0.01$	Supported (pilot)

H1 is supported at pilot scale: CSMM maturity is positively associated with value ($b = +0.642$, $p < 0.001$; adjusted $R^2 = 0.581$), a share of explained variance that should shrink in a larger, more heterogeneous sample. H2 is partially supported: the CSMM $\times$ TRI interaction is negative and significant (-0.315 , $p < 0.015$), but the direct TRI effect (-0.210 , $p = 0.082$) is not, and is not reported as an established main effect. **Because Equation (1) implies a negative interaction by construction (Proposition 1), the interaction is informative only if the outcome is the observed value measure (see the note in Section 3).** For H3, the compressive form fits with $f = 5.20\%$ and $k = 0.412$ ($p < 0.01$) around a baseline cost of equity of 8.50% that was fixed as an industry-benchmark assumption, not estimated; this is a first pilot calibration of Equations (3)–(4), not an independent estimate. **Equation (3) defines f as a dimensionless beta ratio, whereas the reported 5.20% is expressed against a cost of equity, and Section 5.1 shows f to be weakly identified.**

5.2.5 Qualitative Barriers

Open-ended responses converged on three barriers: multi-cloud telemetry fragmentation and incompatible Scope 3 carbon estimates across providers (45%), architectural paralysis from overlapping CSRD/DORA/AI Act timelines (35%), and hyperscaler lock-in limiting the cross-border relocation that SSCTB and Sovereign Resilience presuppose (20%). **These map onto the CE/RA and SR/VDR items flagged in the Delphi phase and onto the SSCTB preconditions (Section 3.1), giving qualitative texture to the quantitative associations.**

5.3 Multi-Criterion Evaluation of the Evidence

The pilot is evaluated against ten criteria, each stated with its test, the present evidence, its status and the study still required (Table 7). **The evaluation separates formal validity, which the paper establishes analytically, from empirical validity, which the pilot supports only directionally.**

Table 7. Multi-criterion evaluation of the framework

Criterion	Test applied	Present evidence	Status	Required next
Mathematical validity	Analytic proof (Propositions 1–2)	SCV $\in [0,100]$; $\beta_{post} > 0$ and strictly decreasing in SCV	Established analytically	Empirical calibration of f and k
Content validity	Delphi, 1–7 ratings	5 of 6 groups retained in Round 1; SI	Supported at pilot scale	Independent panel; response

Criterion	Test applied	Present evidence	Status	Required next
		revised, consensus in Round 2		rate per round
Internal consistency	Cronbach's alpha	0.76–0.91; composite 0.87	Directional	Alpha with confidence intervals at $n \geq 100$
Structural validity	CFA, two threshold tiers	Conventional tier met; strict tier met on SRMR only; item-level model not identified at $n = 20$	Not established	Split-half EFA/CFA at $n \geq 100$
Discriminant validity	Not yet tested	None	Outstanding	Formal discriminant-validity tests at $n \geq 100$
Criterion validity	Regression, H1–H3	H1 supported; H2 partial; H3 first calibration	Directional	Observed outcomes and market beta data
Robustness to weights	Analytic sensitivity (Figure 2); sector weights proposed	Structural sensitivity shown; empirical sensitivity untested	Partial	Re-estimate under equal and sector-stratified weights
Common-method bias	Harman's test	Not reported for the pilot	Outstanding	Report result and procedural remedies
Independence and replicability	Cohort, administrator, data access	Same cohort; developer-administered; data not yet deposited	Not established	Independent replication; open deposit
Generalisability	Sampling frame	Convenience sample, $n = 20$, six member states	Not established	Stratified sample, $n \geq 100$

A prospective sensitivity analysis, computed for this paper, quantifies the power constraint. For a single-coefficient F test at $\alpha = 0.05$ and power 0.80 in a model with five predictors (CSMM, TRI, their interaction and two controls), $n = 20$ can detect only effects of $f^2 \approx 0.45$ (partial $R^2 \approx 0.31$), whereas $n = 100$ detects $f^2 \approx 0.08$ (partial $R^2 \approx 0.07$). The pilot's adjusted R^2 of 0.581 corresponds to $f^2 \approx 1.39$. A pilot that can reach

significance only for very large effects will tend to overstate those that it does detect, so the pilot coefficients should be treated as upper-bound indications and the $n \geq 100$ design as the minimum credible test. Additional sector dummies would reduce power slightly.

Taken together, the architecture is a pilot-tested design contribution: its formal properties are proven, its measurement is internally consistent at pilot scale, and its empirical claims await the independent study specified in Section 8.

6. DISCUSSION

6.1 Theoretical Implications

The corrected formulations show that a common critique of practitioner-originated maturity frameworks, that their formal models fail basic dimensional and boundary analysis, can be answered without abandoning the underlying proposition (Propositions 1–2). **A design principle follows: a maturity-to-value mapping used for financial inference should be bounded at both ends and monotone, with each parameter identified by data inside its domain.** The weak identification of f (Section 5.1) shows what is lost otherwise. At pilot scale, CSMM predicted value and the fitted compression parameter $k = 0.412$ gives Equation (3) its first real-data anchor, though a sample of 20 cannot bear the full theoretical weight of that claim. The pattern in H2, a significant interaction beside a non-significant direct TRI effect, is consistent with the dynamic-capabilities reading that risk acts on the conversion of maturity into value rather than on value directly (Section 2.1); it is a conjecture for the full-scale study, not a finding.

6.2 Managerial Implications for EU Enterprises

The framework offers finance, security and sustainability functions a common scoring language. Because TRI enters both the value composite and the cost-of-capital adjustment, it gives enterprises a structured way to argue, once validated, that resilience and sovereignty investment is a lever on value and financing cost, not purely a cost centre. The barriers executives named (Section 5.2.5) map onto the SSCTB preconditions, suggesting that the structural layer, not only the scoring layer, is where practitioners currently feel most friction.

6.3 Policy Implications

The mapping in Table 1 suggests that CSRD/ESRS, NIS2/DORA and AI Act compliance could in principle be assessed through a common maturity instrument, potentially reducing the compliance burden of cloud-dependent SMEs. The panel's recommendation to weight Regulatory Assurance and Sovereign Resilience more heavily in regulated sectors is a small piece of supporting evidence, which, like the rest of the pilot, awaits full-scale confirmation before it informs policy design.

7. LIMITATIONS AND THREATS TO VALIDITY

- **Model status.** Weights and the parameters f and k are calibrated only at pilot scale; f is weakly identified (Section 5.1); the 8.50% baseline cost of equity is a fixed assumption; and the parametrisation of f requires reconciliation (Section 5.2.4).
- **Measurement and outcome definition.** Perceived value and model-implied SCV must not be conflated (Section 3). Self-report and common-method variance affect any survey collecting predictors and outcomes from one respondent, and are compounded

by the shared pilot sample. Item-level CFA is not identified at $n = 20$.

- **Sampling and independence.** The pilot is a convenience sample of 20 senior EU-based decision-makers. The Delphi panel and survey draw on the same individuals, so agreement between them is not independent triangulation. The pilot was administered by the framework's developer, who also originated its constructs (see Conflicts of Interest).
- **Inference.** The direct TRI effect is not significant, so H2 is partially supported. A cross-sectional design cannot support maturity-transition claims. Endogeneity and reverse causality are possible, since well-resourced organisations may invest in maturity because they are already valuable. The illustrative profiles have no counterfactual and are hypothetical.
- **Reporting.** The pilot dataset is not yet deposited (see Data Availability), and any claim of pre-registration requires a registry identifier.

8. CONCLUSIONS, RECOMMENDATIONS AND FUTURE WORK

This paper has restructured a practitioner-originated cloud sustainability framework into an internally consistent, mathematically bounded architecture. It corrected a scale mismatch and an unbounded cost-of-capital function, proved the boundary and monotonicity properties of the corrected model, specified fieldable instruments, and reported a pilot ($n = 20$) that offers directional support for H1, partial support for H2 and a first calibration for H3. The pilot's small, shared, developer-administered sample means its statistics are directional, and the direct TRI-to-value effect did not reach significance.

8.1 Recommendations

Researchers: run the independent $n \geq 100$ study with a sample separate from the Delphi panel, report all planned statistics regardless of outcome, and re-test the direct TRI effect. *Enterprises:* use SSCTB as a readiness checklist before CSMM self-scoring; multi-cloud telemetry and vendor lock-in are the most immediate gaps. *Policymakers:* treat the regulation-to-framework mapping and the sector-weighting recommendation as hypotheses for independent research, particularly on SME compliance burden. *Reviewers and editors:* check any future reported alpha, loading, fit index or coefficient against a disclosed, replicable dataset.

8.2 Future work

(i) The full, independent $n \geq 100$ survey, reported as specified in Section 4.4 and comparable with the pilot estimates; (ii) a longitudinal panel across at least two CSMM assessment cycles 18–24 months apart; (iii) an independent replication by researchers not involved in the framework's design; and (iv) re-estimation of f and k against disclosed listed-company beta data with the parametrisation reconciled. Only then would a claim of full empirical validation be defensible.

9. ACKNOWLEDGMENT

The author thanks the JAAI referee for constructive comments that improved the manuscript.

Conflicts of Interest. The framework constructs (SSCTB, CSMM, SCVF, TRI) originate in the author's own practitioner

work. On the basis for the recommendation that validation be replicated by independent researchers.

Data Availability. The pilot dataset ($n = 20$: Delphi ratings and survey responses underlying Section 5.2 and Tables 4–6) exists as an internal analytical report and has not been deposited. On completion of the full-scale protocol, the pilot data, anonymised full-scale responses, reliability and validity statistics and analysis code will be deposited in an open repository (for example OSF or Zenodo) and the DOI cited thereafter.

Ethics. The pilot Delphi panel and survey ($n = 20$) were administered directly by the framework's developer. The empirical work was an expert-consensus exercise on the content validity, relevance and clarity of the framework constructs, with a corresponding practitioner survey; participants were experienced cloud, enterprise-architecture, cybersecurity and digital-governance professionals contributing in a professional capacity. No vulnerable populations, clinical intervention, deception, behavioural experimentation or special-category data under Regulation (EU) 2016/679 were involved. Participation was voluntary, responses were anonymised before analysis and results are reported in aggregate only. The study was assessed against institutional research-ethics guidance for computing and information-systems research and determined to present minimal ethical risk; formal institutional approval was accordingly not required for the pilot. Formal ethics review will be sought where the larger sample warrants it.

Author Contributions. Justice Opara-Martins: conceptualisation, methodology, formal analysis, investigation, writing and supervision.

Supplementary Material. S1: Delphi Round 1 questionnaire (formerly Appendix A). S2: practitioner survey instrument (formerly Appendix B). Both are unchanged from the submitted manuscript apart from the corrections noted in Section 4.2. template.

10. REFERENCES

- [1] 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, 16.12.2022.
- [2] European Commission. (2023). Commission Delegated Regulation (EU) 2023/2772 of 31 July 2023 supplementing Directive 2013/34/EU of the European Parliament and of the Council as regards sustainability reporting standards (ESRS). Official Journal of the European Union, 22.12.2023.
- [3] Directive (EU) 2022/2555 of the European Parliament and of the Council of 14 December 2022 on measures for a high common level of cybersecurity across the Union, amending Regulation (EU) No 910/2014 and Directive (EU) 2018/1972, and repealing Directive (EU) 2016/1148 (NIS2). Official Journal of the European Union, L 333, 27.12.2022.
- [4] 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, 27.12.2022.

- [5] Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 laying down harmonised rules on artificial intelligence (AI Act). Official Journal of the European Union, 12.7.2024.
- [6] Vial, G. (2019). Understanding digital transformation: A review and a research agenda. *Journal of Strategic Information Systems*, 28(2), 118–144. <https://doi.org/10.1016/j.jsis.2019.01.003>
- [7] Teece, D. J. (2018). Business models and dynamic capabilities. *Long Range Planning*, 51(1), 40–49.
- [8] Barney, J. B. (1991). Firm resources and sustained competitive advantage. *Journal of Management*, 17(1), 99–120. <https://doi.org/10.1177/014920639101700108>
- [9] International Organization for Standardization. (2015). ISO 14001:2015 – Environmental management systems – Requirements with guidance for use.
- [10] International Organization for Standardization. (2022). ISO/IEC 27001:2022 – Information security, cybersecurity and privacy protection – Information security management systems – Requirements.
- [11] International Organization for Standardization. (2023). ISO/IEC 42001:2023 – Information technology – Artificial intelligence – Management system.
- [12] Hsu, C.-C., & Sandford, B. A. (2007). The Delphi technique: Making sense of consensus. *Practical Assessment, Research & Evaluation*, 12(10). <https://doi.org/10.7275/pdz9-th90>
- [13] Cronbach, L. J. (1951). Coefficient alpha and the internal structure of tests. *Psychometrika*, 16(3), 297–334. <https://doi.org/10.1007/BF02310555>
- [14] Hu, L., & Bentler, P. M. (1999). Cutoff criteria for fit indexes in covariance structure analysis: Conventional criteria versus new alternatives. *Structural Equation Modeling*, 6(1), 1–55. <https://doi.org/10.1080/10705519909540118>
- [15] Podsakoff, P. M., MacKenzie, S. B., Lee, J.-Y., & Podsakoff, N. P. (2003). Common method biases in
- Carbon Efficiency (CE)

behavioral research: A critical review of the literature and recommended remedies. *Journal of Applied Psychology*, 88(5), 879–903. <https://doi.org/10.1037/0021-9010.88.5.879>

APPENDIX A. Delphi Panel - Round 1 Questionnaire

This instrument operationalises the Delphi protocol specified in Section 4.2 and should be read alongside it.

Data-collection instrument. A pilot administration of this instrument (n = 20) has been completed and is reported in Section 5.2; the item wording below is the version used for that pilot and reflects the Round 2 revision to the Sustainable Innovation items noted in Section 5.2.1. Results (IQR, % agreement per item) are reported in Table 4.

A.1 Instructions for Panellists

- For each candidate item, rate Relevance, Clarity, and Completeness on a 1-7 scale.
- Use the comment line under each construct to suggest rewording or flag missing items.
- Responses are confidential; only aggregated, anonymised results are reported or shared with other panellists in Round 2.
- Estimated time: 25-35 minutes. Consensus criterion: IQR ≤ 1 and $\geq 75\%$ of panellists rating Relevance ≥ 5 .

A.2 Panellist Background

- Primary area of expertise: ☐ Cloud Architecture ☐ Sustainability/ESG ☐ EU Regulatory Policy ☐ Academic Research ☐ Other
- Years of relevant experience: ☐ <5 ☐ 5-10 ☐ 11-20 ☐ 20+
- Conflict of interest to declare: ☐ None ☐ Yes - describe: _____

A.3 CSMM Construct Item Evaluation

- Carbon Efficiency (CE)

Candidate item	Relevance (1-7)	Clarity (1-7)	Completeness (1-7)
Our cloud workloads are scheduled with carbon intensity as an explicit optimisation input.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
We measure the carbon footprint of individual workloads, not only total data-centre energy use.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
Renewable-energy availability influences where and when we run non-time-critical workloads.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
We can report workload-level carbon data to auditors on request without manual reconstruction.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]

- Regulatory Assurance (RA)

Candidate item	Relevance (1-7)	Clarity (1-7)	Completeness (1-7)
We can demonstrate CSRD/ESRS-aligned reporting for cloud-hosted operations without manual reconciliation.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
Our cloud architecture has a named owner accountable for sustainability-regulation compliance.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
We routinely map new EU regulatory requirements (e.g. AI Act obligations) onto our cloud architecture.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]

Candidate item	Relevance (1-7)	Clarity (1-7)	Completeness (1-7)
Our compliance evidence for cloud operations is produced automatically rather than compiled ad hoc.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]

- Sovereign Resilience (SR)

Candidate item	Relevance (1-7)	Clarity (1-7)	Completeness (1-7)
We have documented control over the jurisdiction and access conditions of our critical data.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
We could relocate critical workloads to another jurisdiction within a defined recovery time objective.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
Our contracts with cloud vendors specify data-sovereignty guarantees, not only availability SLAs.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
We have tested (not only documented) our data-sovereignty controls in the last 12 months.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]

- Operational Optimisation (OO)

Candidate item	Relevance (1-7)	Clarity (1-7)	Completeness (1-7)
AI-assisted tooling actively rebalances our workloads for cost and performance.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
We use predictive (not only reactive) alerting for capacity and performance issues.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
Our infrastructure-as-code practices allow architecture changes to be tested before production deployment.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
Optimisation decisions are based on integrated cost/carbon/performance data, not cost alone.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]

- Sustainable Innovation (SI)

Candidate item	Relevance (1-7)	Clarity (1-7)	Completeness (1-7)
Sustainability criteria are formally scored in our cloud vendor/architecture selection process.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
We pilot emerging sustainable-cloud practices ahead of mandatory adoption.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
Sustainability outcomes are part of how our cloud/IT leadership is evaluated.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
We have a defined roadmap for sustainability-linked cloud investment over the next 24 months.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]

A.4 TRI Item Evaluation

Items marked (reverse-scored) are worded in the protective direction and will be reverse-coded during analysis.

- Regulatory Risk (RR)

Candidate item	Relevance (1-7)	Clarity (1-7)	Completeness (1-7)
We are currently uncertain how a specific EU regulation applies to part of our cloud estate.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]

Candidate item	Relevance (1-7)	Clarity (1-7)	Completeness (1-7)
A regulatory audit of our cloud operations would likely surface undocumented gaps.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]

- Cyber Risk (CR)

Candidate item	Relevance (1-7)	Clarity (1-7)	Completeness (1-7)
Our cloud environment has unresolved findings from our most recent security assessment.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
We have a tested incident-response plan specific to cloud-hosted critical services. (reverse-scored)	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]

- Vendor-Dependency Risk (VDR)

Candidate item	Relevance (1-7)	Clarity (1-7)	Completeness (1-7)
A single cloud vendor's exit would materially disrupt our operations within 90 days.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
We could migrate our most critical workload to another provider within 6 months if required. (reverse-scored)	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]

- Operational-Disruption Risk (ODR)

Candidate item	Relevance (1-7)	Clarity (1-7)	Completeness (1-7)
Unplanned downtime in the last 12 months has materially affected business operations.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
We have redundant architecture for our most business-critical cloud workload. (reverse-scored)	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]

- Sovereignty-Compliance Risk (SCR)

Candidate item	Relevance (1-7)	Clarity (1-7)	Completeness (1-7)
We are not fully confident our current data residency matches our contractual/regulatory commitments.	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]
We have documented evidence of where our regulated data physically resides at all times. (reverse-scored)	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]	[1 2 3 4 5 6 7]

A.5 Construct-Level and Architecture Feedback

- SCVF weighting: should all five CSMM constructs be weighted equally, or should some (e.g. Regulatory Assurance) carry higher default weight for EU enterprises?
- SSCTB structural checklist: are the four preconditions necessary and sufficient? What would you add, remove, or merge?
- Missing constructs: is there a dimension of cloud sustainability, resilience, or value not captured by CSMM, TRI, or SCVF as defined?

APPENDIX B. PRACTITIONER SURVEY INSTRUMENT

This instrument operationalises the quantitative survey protocol specified in Section 4.3 and should be read alongside it

A pilot administration of this instrument (n = 20, the same cohort as Appendix A) has been completed and is reported in Section 5.2; the results in Tables 5-6 are pilot-scale only. The instrument below should now be fielded at full scale (target n >= 100) with a sample independent of the Delphi panel, per Section 8's recommendation.

B.1 Introduction and Consent

Approximately 12-15 minutes. Participation is voluntary; responses are confidential and reported only in aggregate. In

line with GDPR, no personal data beyond professional role and organisation type is collected; responses are stored securely; participants may withdraw before submission without providing a reason.

B.2 Screening and Organisation Profile

- Primary role: ☐ Cloud/IT Architecture ☐ IT/Infrastructure Leadership ☐ Sustainability/ESG ☐ Risk/Compliance ☐ Finance ☐ Other
- Sector: ☐ Financial Services ☐ Manufacturing ☐ Public Sector ☐ Retail/Consumer ☐ Technology ☐ Healthcare ☐ Other
- Carbon Efficiency (CE)

- Organisation size: ☐ <250 (SME) ☐ 250-999 ☐ 1,000-4,999 ☐ 5,000+
- Primary EU member state of operation: _____
- Operates workloads in public cloud infrastructure: ☐ Yes ☐ No (if No, survey ends)

B.3 CSMM Items (1=Strongly disagree ...

5=Strongly agree)

Item	1=Strongly disagree ... 5=Strongly agree
CE1. Our cloud workloads are scheduled with carbon intensity as an explicit optimisation input.	[1 2 3 4 5]
CE2. We measure the carbon footprint of individual workloads, not only total data-centre energy use.	[1 2 3 4 5]
CE3. Renewable-energy availability influences where and when we run non-time-critical workloads.	[1 2 3 4 5]
CE4. We can report workload-level carbon data to auditors on request without manual reconstruction.	[1 2 3 4 5]

- Regulatory Assurance (RA)

Item	1=Strongly disagree ... 5=Strongly agree
RA1. We can demonstrate CSRD/ESRS-aligned reporting for cloud-hosted operations without manual reconciliation.	[1 2 3 4 5]
RA2. Our cloud architecture has a named owner accountable for sustainability-regulation compliance.	[1 2 3 4 5]
RA3. We routinely map new EU regulatory requirements (e.g. AI Act obligations) onto our cloud architecture.	[1 2 3 4 5]
RA4. Our compliance evidence for cloud operations is produced automatically rather than compiled ad hoc.	[1 2 3 4 5]

- Sovereign Resilience (SR)

Item	1=Strongly disagree ... 5=Strongly agree
SR1. We have documented control over the jurisdiction and access conditions of our critical data.	[1 2 3 4 5]
SR2. We could relocate critical workloads to another jurisdiction within a defined recovery time objective.	[1 2 3 4 5]
SR3. Our contracts with cloud vendors specify data-sovereignty guarantees, not only availability SLAs.	[1 2 3 4 5]
SR4. We have tested (not only documented) our data-sovereignty controls in the last 12 months.	[1 2 3 4 5]

- Operational Optimisation (OO)

Item	1=Strongly disagree ... 5=Strongly agree
OO1. AI-assisted tooling actively rebalances our workloads for cost and performance.	[1 2 3 4 5]
OO2. We use predictive (not only reactive) alerting for capacity and performance issues.	[1 2 3 4 5]
OO3. Our infrastructure-as-code practices allow architecture changes to be tested before production deployment.	[1 2 3 4 5]
OO4. Optimisation decisions are based on integrated cost/carbon/performance data, not cost alone.	[1 2 3 4 5]

- Sustainable Innovation (SI)

Item	1=Strongly disagree ... 5=Strongly agree
SI1. Sustainability criteria are formally scored in our cloud vendor/architecture selection process.	[1 2 3 4 5]
SI2. We pilot emerging sustainable-cloud practices ahead of mandatory adoption.	[1 2 3 4 5]
SI3. Sustainability outcomes are part of how our cloud/IT leadership is evaluated.	[1 2 3 4 5]
SI4. We have a defined roadmap for sustainability-linked cloud investment over the next 24 months.	[1 2 3 4 5]

B.4 TRI Items

Items marked [R] are reverse-scored during analysis; respond to the statement exactly as written.

- Regulatory Risk (RR)

Item	1=Strongly disagree ... 5=Strongly agree
RR1. We are currently uncertain how a specific EU regulation applies to part of our cloud estate.	[1 2 3 4 5]
RR2. A regulatory audit of our cloud operations would likely surface undocumented gaps.	[1 2 3 4 5]

- Cyber Risk (CR)

Item	1=Strongly disagree ... 5=Strongly agree
CR1. Our cloud environment has unresolved findings from our most recent security assessment.	[1 2 3 4 5]
CR2. We have a tested incident-response plan specific to cloud-hosted critical services. [R]	[1 2 3 4 5]

- Vendor-Dependency Risk (VDR)

Item	1=Strongly disagree ... 5=Strongly agree
VDR1. A single cloud vendor's exit would materially disrupt our operations within 90 days.	[1 2 3 4 5]
VDR2. We could migrate our most critical workload to another provider within 6 months if required. [R]	[1 2 3 4 5]

- Operational-Disruption Risk (ODR)

Item	1=Strongly disagree ... 5=Strongly agree
ODR1. Unplanned downtime in the last 12 months has materially affected business operations.	[1 2 3 4 5]
ODR2. We have redundant architecture for our most business-critical cloud workload. [R]	[1 2 3 4 5]

- Sovereignty-Compliance Risk (SCR)

Item	1=Strongly disagree ... 5=Strongly agree
SCR1. We are not fully confident our current data residency matches our contractual/regulatory commitments.	[1 2 3 4 5]
SCR2. We have documented evidence of where our regulated data physically resides at all times. [R]	[1 2 3 4 5]

B.5 Outcomes (Sustainable Cloud Value)

- Perceived Value Outcomes

Item	1=Strongly disagree ... 5=Strongly agree
SCV1. Our cloud investment decisions are made with a clear view of their financial return.	[1 2 3 4 5]
SCV2. Sustainability-related cloud investment has been linked, in practice, to a measurable reduction in operating cost.	[1 2 3 4 5]
SCV3. Our organisation's ESG or sustainability rating has been positively influenced by our cloud architecture choices.	[1 2 3 4 5]
SCV4. Investors, auditors, or the board have referenced our cloud sustainability practices favourably in the last 12 months.	[1 2 3 4 5]

B.6 Open Response (Optional)

- What is the single biggest barrier to advancing your organisation's cloud sustainability maturity?
- Is there anything about your organisation's approach to cloud sustainability, risk, or value that this survey did not capture?