

Blockchain or Not? An Auditable Decision Framework for Organisations

Original

Blockchain or Not? An Auditable Decision Framework for Organisations / Meneguzzo, S., Leto, S., Romani, N., Gatteschi, V., Favenza, A., Schifanella, C.. - (2026), pp. 238-247. (GoodIT '26: International Conference on Information Technology for Social Good Pisa (ITA) September 2 - 4, 2026) [10.1145/3794786.3830760].

Availability:

This version is available at: 11583/3016371 since: 2026-10-07T05:13:51Z

Publisher:

Association for Computing Machinery (ACM)

Published

DOI:10.1145/3794786.3830760

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Blockchain or Not? An Auditable Decision Framework for Organisations

Silvio Meneguzzo*

Computer Science
University of Turin
Torino, Torino, Italy
FCC - Blockchain
Links Foundation
Torino, Torino, Italy
meneguzzosilvio@gmail.com

Stefano Leto

FCC - Blockchain
LINKS Foundation
Torino, Torino, Italy
stefano.letto@linksfoundation.com

Noemi Romani

Computer and Control Engineering
Politecnico di Torino
Torino, Torino, Italy
noemi.romani@polito.it

Valentina Gatteschi

Computer and Control Engineering
Politecnico di Torino
Torino, Torino, Italy
valentina.gatteschi@polito.it

Alfredo Favenza

FCC - Blockchain
Links Foundation
Torino, Torino, Italy
alfredo.favenza@linksfoundation.com

Claudio Schifanella

Computer Science
University of Turin
Torino, Torino, Italy
claudio.schifanella@unito.it

Abstract

Determining when *not* to adopt blockchain or distributed ledger technology (DLT) remains difficult, especially in public-sector and social-good initiatives where unjustified experimentation can waste scarce resources, increase compliance burden, and weaken institutional accountability. This paper presents a compact, evidence-based phase-gate framework for auditable blockchain adoption decisions. The framework structures the decision as a sequence of bounded phases and must-meet gates, each producing minimal evidence artefacts and explicit go/no-go outcomes. When a decisive condition is not satisfied, the process exits towards a documented non-DLT alternative rather than compensating the failure through later scoring.

The procedure separates necessity from implementability. First, a necessity triage verifies lawful immutable processing, independent writers operating under limited trust, and a genuine requirement for cross-boundary immutability. Only then does the framework assess permission and disclosure fit, IT integration, smart-contract maintainability, performance and finality under stated Service Level Objectives (SLOs), and economic sustainability. Focused subflows support legal and compliance analysis, integration readiness, and performance assessment. Derived from a structured review and comparative mapping, the framework is appraised analytically and through a formative walk-through with first-time users at a commercial engineering-software vendor. The results indicate clearer rejection pathways, improved decision transparency, and a more auditable route to adoption when blockchain is warranted.

*Corresponding author.



This work is licensed under a Creative Commons Attribution 4.0 International License. *GoodIT '26, Pisa, Italy*

© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2483-1/26/09
<https://doi.org/10.1145/3794786.3830760>

CCS Concepts

• **Information systems** → **Information systems applications**; • **Software and its engineering** → **Software notations and tools**; • **Social and professional topics** → **Computing / technology policy**.

Keywords

Blockchain adoption, Distributed ledger technology, Phase-gate decision framework, Go/no-go gates, Permissioned vs. permissionless, Governance and compliance, Systems integration, Social good

ACM Reference Format:

Silvio Meneguzzo, Stefano Leto, Noemi Romani, Valentina Gatteschi, Alfredo Favenza, and Claudio Schifanella. 2026. Blockchain or Not? An Auditable Decision Framework for Organisations. In *International Conference on Information Technology for Social Good (GoodIT '26)*, September 02–04, 2026, Pisa, Italy. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3794786.3830760>

1 Introduction

Public administrations, civic-participation platforms, and social-economy initiatives increasingly pilot blockchain for transparency and trust, yet uncritical adoption wastes scarce public resources, complicates compliance, and erodes the very accountability that decentralised ledgers are meant to support. Deciding *when not* to adopt is therefore itself a social good: it protects budgets, citizens' data, and institutional legitimacy.¹

Prior work [1, 5, 8, 13, 29, 32] converges on recurrent decision axes (trust distribution, shared writes, performance/finality, privacy, governance/legal fit, economics, integration), yet organisations lack a compact, operational template turning these axes into defensible decisions usable by mixed technical-managerial teams. Following Cooper's Stage-Gate tradition [9], we treat *phases* as bounded work segments producing minimal evidence artefacts and *gates* as discrete go/no-go checkpoints; we adopt his *must-meet* (non-compensatory) posture, so that a single failure halts the blockchain

¹Extended notation, related-work synthesis, and contribution rationale are released in the public Figshare companion.

path and triggers a documented off-ramp, with no offsetting from other criteria. Early gates form a *necessity triage* (lawful processing; genuinely independent writers; acceptability of a continuously online trusted third party (TTP) and cross-boundary immutability); only upon passing these gates does the process advance to *implementability* gates (integration feasibility, performance and finality against SLOs, economic sustainability). This ordering avoids compensating decisive disqualifiers and grounds outcomes in concrete artefacts, responding to BPM-literature calls to connect conceptual blockchain guidance to deployable, lifecycle-aware processes [24].

Notation. *Phases* (P0–P10) assemble inputs and produce minimal evidence artefacts; *gates* (G0–G9) test those artefacts and issue go/no-go decisions, with mapping $P_i \rightarrow G_i$ for $i \in \{0, \dots, 7\}$ (P8 feeds G7/G8; P9→G8; P10→G9). The full catalogue and iconography are deferred to the public artefact (see footnote 1).

Contributions. Against the four documented gaps (non-auditable reject-first stops; the necessity being entangled with implementability; uneven off-ramps; and triage-heavy scoring [5, 8, 13, 29, 32]), the framework contributes: (C1) *reject-first logic made auditable* via discrete go/no-go checkpoints with minimal evidence artefacts; (C2) *compliance and integration as first-class decisions*, with dedicated subflows (inputs, checks, reject criteria); (C3) *necessity separated from implementability*, ordering lawful immutability and limited-trust multi-party writes ahead of platform choices; and (C4) *lightweight post-gate assessment*, applying structured scoring only after the hard checkpoints. Practical uptake is appraised through a formative walk-through with first-time users at a commercial building-energy/HVAC software vendor whose toolchain spans certification, compliance, and OpenBIM/IFC exchange² (Section 4.3). A Figshare companion deposit [4] archives the full SLR protocol (search strings, PRISMA accounting, exclusion log), the extended phase/gate notation, the high-resolution flowcharts, and operational-tooling references that are summarised in the body for space. The paper proceeds with §2 (review protocol and related work), §3 (the phase–gate framework), §4 (analytical and formative results), and §5 (conclusions and threats).

2 Definitions and Related Work

2.1 Scope, Definitions, and Review Protocol

This review analyses decision frameworks that determine whether and when to adopt blockchain or distributed ledger technology (DLT) over non-DLT alternatives, paying particular attention to permissionless versus permissioned designs. Three sources were queried using mirrored Boolean predicates: IEEE Xplore, Semantic Scholar, and Elsevier ScienceDirect. IEEE Xplore was used as the operational counting frame for PRISMA accounting, while Semantic Scholar and Elsevier ScienceDirect were used for triangulation, backward and forward citation chaining, and coverage checks rather than being pooled into the denominator.

Studies were included if they: (i) present an explicit adoption framework that can be rejected in favour of non-DLT options [8, 19, 29, 32]; (ii) treat permission models distinctly [5, 13, 15]; and (iii) cover multi-criteria dimensions spanning technical, organisational,

governance, economic, legal, privacy, and integration concerns [1, 6, 23, 28].

Reporting follows the PRISMA 2020 statement [27]. The IEEE Xplore counting frame yielded 1,399 unique records after deduplication, all screened on title and abstract; 1,052 records were excluded at this stage as out of scope. Of the 347 reports sought for retrieval, 41 could not be retrieved because they were paywalled, associated with dead DOI links, available only as conference abstracts, or not accessible to the review team in a language they understood, leaving 306 reports assessed for eligibility. Of these, 265 were excluded for reasons recorded in the exclusion log (failure to meet inclusion criteria (i)–(iii)), and 41 studies were included in the final review corpus. Full search strings, inclusion and exclusion criteria, the PRISMA 2020 flow diagram, stage-by-stage accounting, the retrieval-failure breakdown, screening predicates, and the exclusion log are deposited as Supplementary Material A–D in the companion Figshare bundle (see Data Availability).

2.2 Synthesis of Frameworks

Decision criteria concentrate along nine recurrent axes (trust/data mutability, shared writes, smart contracts, asset digitisation, performance/finality, privacy, governance/legal, economics, integration) that surface in keystone frameworks [1, 5, 8, 13, 15, 29, 32] and sector overlays [3, 6, 7, 16, 17, 22, 30]. Three artefact families recur: reject-first decision trees with explicit off-ramps (Wüst & Gervais; Chowdhury; Scriber) [8, 29, 32]; control-based assessments (BAF: 5 domains, 18 subdomains, ~100 controls) and when–which–how vademecums [5, 13]; and weighted-scoring or scoping reviews [1, 15, 19, 25, 28]. Table 1 summarises coverage. Consensus is strong on the permissioned/permissionless split and on the principle of off-ramps; legal/regulatory and privacy coverage is uneven outside public-sector studies, and legacy integration is rarely operationalised as a hard gate. Foundational-framework prose, sector-specific syntheses (construction, logistics, healthcare, energy/EV, public sector), and an extended agree/diverge discussion are released in the public artefact (footnote 1).

Adoption outcomes still hinge on operationalising legal/compliance, legacy integration, and systematic alternative triage (the levers addressed by the framework presented next).

3 Research Methodology

Table 2 reads the literature through the framework’s nomenclature, separating *hard go/no-go gates* (non-compensatory: G0; G1–G4; G7–G9) from *operational features/stages*. Symbols encode each reference’s treatment: ✓ explicit hard stop or executable stage; Δ conceptual or scored only; ✗ not treated as a distinct gate/stage.

Quantitatively, across the five references considered, only 2/5 implement an explicit triage on G1–G4 (Wüst & Gervais; Chowdhury); 0/5 elevate integration (G7) to a gate; 0/5 treat SLO/finality as a hard stop (G8) or economics/incentives as a hard stop (G9); and 3/5 discuss off-ramps to non-DLT options without a non-compensatory decision. Overall, the result is a conservative and auditable process that avoids more complex solutions when simpler ones will suffice and that requires concrete integration and governance artefacts before advancing.

²<https://www.buildingsmart.org/standards/bsi-standards/industry-foundation-classes/>

Table 1: Coverage matrix for representative frameworks.

Reference	Perm.	Non-DLT	Multi-crit.	Artefact	Notes
Wüst & Gervais (2018) [32]	✓	✓	✓	Decision tree	Canonical reject-gates; applied cases
Gourisetti et al. (2020) [13]	✓	✓	✓	Scoring/matrix	~100 controls; used in infra/energy
Chowdhury et al. (2018) [8]	✓	✓	✓	Decision tree	Blockchain vs. database emphasis
Belotti et al. (2019) [5]	✓	✓	✓	Survey/vademecum	Consensus taxonomy; platform contrasts
Scriber (2018) [29]	±	✓	✓	Checklist	Synthesis of 23 implementations
Monrat et al. (2022) [25]	✓	✓	✓	Weighted scoring	Six domains; example use case
Hassija et al. (2021) [15]	✓	✓	✓	Criteria list	Includes non-blockchain DLTs
Almeshal & Alhogail (2021) [1]	✓	✓	✓	Scoping review	Maps existing frameworks; identifies gaps
Capocasale et al. (2025) [7]	✓	✓	–	Case/validation	EV supply chain; live production data
Cagigas et al. (2023) [6]	–	–	✓	Evaluation	Public sector; legal/institutional focus

Table 2: Gate/subflow mapping against established frameworks. The table separates *hard go/no-go gates* from *operational features/stages*. For gate definitions, see Table 3; for coarse coverage, see Table 1.

Legend: ✓ explicit hard stop or executable stage; △ conceptual/scored only; ✗ not treated as a distinct gate/stage.

Hard go/no-go gates	Proposed	Wüst & Gervais [32]	BAF [13]	Scriber [29]	Belotti et al. [5]	Chowdhury et al. [8]
G0 Compliance (hard stop; evidence artefacts)	✓	✗	△	✗	△	✗
G1–G4 Necessity triage (shared writes, TTP, immutability)	✓	✓	△	△	△	✓
G7 Integration feasibility (runbooks)	✓	✗	△	✗	✗	✗
G8 Performance/finality under SLOs	✓	✗	△	△	△	✗
G9 Economics/incentives	✓	✗	△	△	△	✗
Operational features / stages						
G5–G6 Permission model (participants)	✓	✓	△	△	△	△
Disclosure (verifiability vs. confidentiality)	✓	✓	△	△	△	△
Legal/Compliance subflow (roles, rights, records)	✓	✗	△	✗	△	✗
IT integration subflow (identity & key custody, observability, SRE runbooks)	✓	✗	△	✗	✗	✗
Performance subflow (workload, finality)	✓	✗	△	✗	△	✗

The remainder of this section formalises the framework as a phase-gate procedure (P0–P10; G0–G9) and outlines the minimal evidence artefacts required at each decision point, from the initial overview to triage, architectural branches, and focused subflows.

3.1 Framework Overview and Phases

The flow stages early problem-framing checks before architectural feasibility, aligned with comparative DLT design analyses [14, 33]. Two indices summarise progress: *Necessity* (G1–G4) and *Implementability* (G7–G9). A decision-maturity pyramid and zoomed band-diagrams of the triage (G0–G4) and the post-triage architectural/implementability checks (G5–G9) are deposited in the Figshare companion (see footnote 1).

Figure 1 renders the procedure as a phase-gate flow (rectangles = phases P0–P10; diamonds = gates G0–G9 with explicit off-ramps). The subflow diagrams (legal/compliance (Figure 2), IT integration (Figure 3), and performance/scalability (Figure 4)) are invoked only when the corresponding checkpoint in the main flow requires deeper analysis; each subflow returns a pass/fail outcome to its parent checkpoint.

Unit of analysis. The procedure is applied once per *decision unit*: a data class or workflow with a homogeneous trust and disclosure profile, rather than once per organisation or system; Figure 1 depicts one such pass. Systems that mix heterogeneous data classes (e.g., organisational identities suited to public verifiability alongside citizen data requiring restricted disclosure) run the procedure once

per unit, each pass yielding its own evidence artefacts and its own off-ramp or architecture. Hybrid and multi-layer designs, such as a private permissioned ledger anchoring state commitments to a public chain, then emerge as *adoption* outcomes whenever distinct passes resolve the G5–G6 permission/disclosure branch differently, provided each resulting layer satisfies the implementability gates G7–G9; the corresponding anchoring commitments are already captured by the on/off-chain plans required at G0 and G4 [20, 21].

3.2 From Triage to Architectural Selection

Gates G0–G4 form a *minimal-need triage*: blockchain is rejected unless: (i) a lawful basis exists for immutable processing, (ii) independent writers under limited trust are present, and (iii) cross-boundary immutability is required, falling back to centralised DB, TTP-managed DB, or shared DB with ACL and tamper-evident logging. Upon passage, the process selects a permission model (P5–P6), verifies integration (G7), defines the scope of smart-contract logic (P8), and tests performance (G8) and economics (G9). To avoid ambiguity between subflows, each specialised analysis returns a named outcome to its parent gate: *is_compliant* for G0, *is_integrable* for G7, and *meets_slos* for G8. The legal and compliance outcome may also assume an *undefined* state, which triggers the time-boxed on-hold procedure defined in Section 3.3.




Figure 1: Blockchain adoption flowchart (phases and decision gates). A high-resolution version is in the Figshare companion.

Table 3: Phase-gate summary with decisions, off-ramps, and minimal evidence artefacts.

Gate/Phase	Decision	Off-ramp (if No)	Evidence (min.)
G0 - Compliance	Compliant?	Stop: not suitable	Jurisdiction map; GDPR roles (Art. 4(7)–(8), 26); lawful basis (Art. 6/9); DPIA (Art. 35); records schedule; on/off-chain placement plan; MiCA/sectoral memos where applicable.
G1 - Minimal need	Pure tokenisation or cross-boundary state persistence?	Stop	Asset/event model; persistence statement (tokenisation vs. stateful).
G2 - Multi-writer	Multiple independent writers?	Prefer centralised DB	Stakeholder-write matrix; authority/RACI; change control.
G3 - TTP acceptance	Trusted third party acceptable?	Prefer TTP-managed DB	TTP service model; SLAs; escrow/exit plan; cost comparison.
G4 - Immutability	Tamper-evident immutability across boundaries?	Shared DB + ACL with append-only logging	Audit boundary; tamper-evidence rationale; on/off-chain commitment plan [20, 21].
G5–G6 - Permission	Participants known/trusted? Confidentiality vs. verifiability?	Branch; if misaligned, re-evaluate	Participant discovery; identity/governance sketch; confidentiality vs. verifiability rationale [5].
G7 - Integration	IT integration feasible?	Stop: not suitable	Interface inventory; observability; runbooks; readiness checklist [18].
P8 - Contracts	Automated logic? Upgradability?	n/a	ABI; access control; upgrade pattern (transparent proxy, UUPS, or Diamond); upgrade authority; rollback path.
G8 - Performance	Adequate under SLOs?	Consider alternative technology	SLOs; latency/finality budgets; headroom assumptions; stress scenarios [2, 11].
G9 - Economics	Costs/resources sustainable?	Select alternative solution	TCO worksheet; fee model; staffing; decommission/rollback plan.

3.3 P0 - Regulatory/Ethical Pre-screen

P0 executes before any other phase and *always* dispatches to the legal/compliance subflow (Fig. 2), which returns a tri-state flag `is_compliant`: `true` sets **G0** to Yes and control returns to the main flow; `false` terminates the blockchain path with a recommendation for a non-DLT alternative; `undefined` places the gate on hold and opens a time-boxed corrective loop: the missing legal evidence is sought within a fixed deadline, after which G0 is re-evaluated; if the deadline lapses unresolved, `is_compliant` defaults to `false` and the off-ramp is taken. The check covers personal data, safety-critical records, and sector-specific obligations, and produces minimal artefacts reused downstream (scoping memo, preliminary records schedule). Legality and record-keeping duties must be evidenced *ex ante*; otherwise, blockchain is deemed not suitable.

3.3.1 How the Subflow Operates (Fig. 2). The subflow performs three jurisdiction-agnostic checks, instantiated here for the EU regime as a worked example; failure at any stage terminates the blockchain path. (i) *Applicable-regime scoping*: where any crypto-asset feature is in scope, its lawfulness is tested against the relevant regime, in the EU, the Markets in Crypto-Assets Regulation (MiCA) for issuers and CASPs, fully applicable since 30 December 2024.³ (ii) *Data-protection scoping* (in the EU, the GDPR): if personal data within Art. 4(1) are processed, data-protection obligations must be satisfiable by design (lawful basis (Art. 6/9), controller/processor and joint-controller roles (Art. 4(7)–(8), Art. 26), DPIA (Art. 35), and reconciliation of data-subject rights (notably the right to erasure under Art. 17) with ledger immutability via on/off-chain placement). The EDPB Guidelines 02/2025 and CNIL guidance converge on privacy-preserving placement (off-chain payloads with on-chain commitments).⁴ (iii) *Sectoral certification/conformity*: where implicated, expert consultation gates progression. When all three pass, **G0** is set to Yes; this dispatch-and-return pattern grounds G0 in *ex-ante* compliance rather than postponed assumptions.

3.4 P1-P4 - Minimal-need Triage

The P1-P4 block makes early off-ramps explicit and auditable, including the fast-track for pure tokenisation and the preference for conventional databases when single-writer authority holds or cross-boundary immutability is absent (see master flowchart, Figure 1). P1 distinguishes pure tokenisation from cross-boundary state persistence; P2 verifies the presence of genuinely independent writers; P3 assesses the acceptability of an always-online trusted third party; P4 checks for tamper-evident immutability across organisational boundaries. Failure at any step routes to the corresponding non-DLT option already documented in Table 3. This conservative approach reduces the risk of costly reversals later on.

3.5 P5 - Participant Discovery and Permission Model

Once necessity has been evidenced, P5 identifies participants and selects the *permission model*. The decision rule at **G5** is direct: when participants are known and confidentiality is required, prefer the

³Regulation (EU) 2023/1114, OJ L 150, 9.6.2023, p. 40, <https://eur-lex.europa.eu/eli/reg/2023/1114/oj>; expanded discussion in the Figshare companion (see footnote 1).

⁴EDPB, *Guidelines 02/2025 on processing of personal data through blockchain technologies*, v. 1.1 (8 April 2025); CNIL, *Blockchain and the GDPR* (29 October 2018).

Legal & Regulatory Compliance (Sub-flow)




Figure 2: G0 - Legal/compliance subflow: regime scoping (e.g., MiCA), GDPR roles and lawful basis, Art. 17 vs. immutability via on/off-chain placement, and sectoral certification, with explicit reject paths. Full narrative in the Figshare companion (footnote 1).

permissioned branch; otherwise consider the permissionless branch, especially where public verifiability or openness are primary goals. The choice establishes the finality and disclosure properties that will be tested at **G8** and shapes identity, onboarding, and key-management practices that will be validated in the integration

subflow (Figure 3). Studies on platform selection with structured criteria and case evidence also emerge at this stage [12].

3.6 P6 - Disclosure vs. Public Verifiability

P6 aligns confidentiality controls with verifiability needs. In the permissioned branch, if strict access controls are required, favour *private permissioned*; otherwise *public permissioned*. In the permissionless branch, if public verifiability is not essential, the problem framing should be revisited. Outputs include a data-placement plan, access-control predicates, and the verifiability rationale, explicitly linking the permission family to finality/disclosure properties later validated at G8 and to integration preconditions assessed at G7. G5 and G6 are presented together because they jointly determine the permission/disclosure profile of the architecture: G5 fixes *who* is allowed to read and write, G6 fixes *what* is publicly verifiable. They are listed as a pair in Table 2 because, among the frameworks reviewed here, none treats them as separable decisions.

3.7 P7 - Integration Readiness

P7 tests whether the proposed ledger can be hosted and operated within the organisation’s technology estate, using the provisional architecture fixed at G5–G6 (Figure 1). The subflow (Figure 3) branches into *on-premises* or *managed-ICT* hosting; the on-premises branch additionally records a *consensus protocol restrained* note when energy policy rules out high-intensity profiles.

Each branch then assembles a minimal integration inventory (identity and key custody (HSM/KMS or threshold/MPC), oracle/bridge trust, vendor-neutral observability, SRE runbooks, and an initial security posture) to demonstrate that the ledger can be *operated* within incumbent constraints. Readiness draws on absorptive-capacity guidance [18]. The subflow returns *is_integrable*: G7 = *Yes* iff true, otherwise the blockchain option is rejected with an auditable rationale. Multi-party supply-chain settings often favour *hybrid* operating models (centralised hosting and incident response, decentralised data and transaction oversight [26]). The full inventory, permissionless-infrastructure considerations (light/full/archive nodes; energy profile; vendor SLAs and exit), and the operational checklist are deferred to the companion.⁵

3.8 P8 - Smart-contract Functionality and Upgradability

P8 determines whether automated logic is needed and, if so, whether it must be upgradable. Where logic is not required, proceed with base ledger functions; where maintenance is needed, adopt an upgradable pattern (transparent proxy, UUPS, or Diamond) with narrowly scoped governance, an explicit *upgrade authority* (multisig/timelock/DAO), and a tested rollback path, compatible with P0 governance and the P7 processes (key custody, change management, runbooks). Outputs include interface specifications (ABI), the upgrade policy, and the operational governance attached to the upgrade authority. ERC-1822/1967/2535 references and OpenZeppelin implementation links are provided in the companion.⁶

⁵Integration subflow companion (full inventory, permissionless-infrastructure notes, vendor links): Figshare, *integration_subflow_full.md*.

⁶See §6 of *integration_subflow_full.md* for ERC-1822 (UUPS), ERC-1967 (proxy-storage slots), ERC-2535 (Diamond), and the OpenZeppelin Upgrades plugins.




Figure 3: G7 - IT integration subflow: hosting branches (on-prem vs. managed ICT), the *consensus protocol restrained* note under energy constraints, and the integration inventory. Any “No” on a critical check sets G7 = No.

3.9 P9 - Performance and Scalability Envelope

P9 tests whether throughput, latency, and finality meet SLOs with realistic headroom. The assessment assumes draft contracts and credible transaction envelopes (read-write mix, burstiness) to avoid naïve extrapolation from vendor claims [11]. The subflow in Figure 4 derives SLOs, fixes the finality model implied by the permission choice, and evaluates capacity, headroom, back-pressure, and failure modes. Where deterministic finality is required (typical of permissioned platforms running Byzantine-fault-tolerant ordering such as PBFT/IBFT/SmartBFT or HotStuff-family protocols, and of crash-fault-tolerant Raft when only non-Byzantine failures are in scope), the latency budget is set accordingly [2, 10]. The subflow exposes a Boolean *meets_slos* outcome to the main flow: G8 is positive only when *meets_slos*=true, i.e. SLOs hold under stress and failure scenarios. Otherwise, the architecture is reconsidered, including multi-layer fallbacks (e.g., moving throughput-critical workload to a Layer-2 construction anchored to the selected ledger) or alternative technologies. Recent work recommends workload-specific SLOs and modular testbeds [31], while comparative studies of permissioned platforms help bound expectations on latency/finality [10]. Artefacts include a capacity worksheet, a finality/latency budget, and stress assumptions.

3.10 P10 - Economic Sustainability

P10 closes the loop by assessing steady-state and stress-state costs, as well as incentive alignment. If resources and expenses are unsustainable or incentives misaligned, G9 is negative and a different solution is selected. Otherwise, the blockchain decision is finalised. The outputs comprise a cost-benefit analysis and an incentive-compatibility memo, which are consistent with the integration and performance evidence gathered earlier.




Figure 4: P9/G8 - Performance and scalability subflow: workload characterisation (read/write mix, burstiness), SLOs, finality model, capacity headroom, back-pressure and failure modes, stress patterns; includes go/no-go checks for envelope violations.

3.11 Must-meet (Non-compensatory) Gates: a Consolidated View

The framework’s termination behaviour can be summarised by a single must-meet rule: gates G0–G4 and G7–G9 are non-compensatory, meaning any failure will terminate the blockchain path, and there is no alternative that can offset that failure. Two substantive points sharpen the off-ramps already listed in Table 3. Firstly, at G4, tamper-evident append-only logging via a Merkle-tree-backed transparency log can already meet cross-boundary immutability requirements [20, 21]. Secondly, at G8, asserted latency/finality budgets or vendor claims are insufficient without empirical support against the target workload [11]. The next section assesses whether this stopping behaviour holds up under analytical comparison, a formative walk-through, and a qualitative gate ablation.

4 Results and Discussion

This section interprets the proposed framework against established baselines and the underlying decision problem. The focus is three-fold: (i) whether the design behaves as a conservative, auditable filter (*when not to use blockchain*); (ii) what each hard checkpoint contributes to that behaviour; and (iii) what limitations and practical implications follow in regulated and managerial settings.

4.1 Scope and Evidence Base

The appraisal is formative rather than summative. It is intended to test analytical coherence, practical intelligibility, and decisional sufficiency, not to provide a statistically generalisable validation of the framework. Three complementary lines of evidence are therefore used: an analytical assessment against canonical decision axes (Section 4.2); a single-session formative walk-through with first-time users (Section 4.3); and a qualitative gate ablation (Section 4.4).

The walk-through was designed as an artefact-centred plausibility test. It examined whether non-specialist users could follow the gate logic, identify the required evidence artefacts, and reach documented off-ramp decisions where appropriate. It was not designed as a controlled usability experiment. Consequently, the findings support claims about practical intelligibility and decisional sufficiency within the observed setting, while broader claims about effectiveness require multi-site replication with recorded think-aloud protocols, independent coding, and inter-rater agreement analysis. The proposed replication design is provided in Supplementary Material E.

4.2 Analytical Assessment

Analytically, the framework makes *reject-first behaviour* explicit: non-compensatory checkpoints (G0, G1–G4, G7–G9) coupled with documented off-ramps prevent compensating decisive disqualifiers and cohere with canonical trees that favour databases or notarised logs when multi-party mistrust or shared-write conditions do not hold [8, 29, 32]. Coherence with established criteria is maintained: the necessity triage and subsequent permission branching mirror decision axes synthesised in control-based assessments and surveys [1, 5, 13]. As summarised earlier (Table 2), compliance (G0), legacy integration (G7), and performance/finality (G8) are elevated to hard go/no-go checkpoints supported by minimal artefacts, while economics/incentives (G9) is treated as an explicit decision rather than a post hoc consideration.

The ordering answers the core question, *when not to use blockchain*, before any platform choice is made. It operationalises the literature’s preference for simpler architectures under single-writer authority, acceptance of a continuously online trusted third party, or, in the absence of cross-boundary immutability, constraining later stages to cases where necessity has been demonstrated. Conceptually, the design shifts emphasis from fully compensatory scoring to a constraint-first approach: legality, minimal need, and absorptive capacity act as binding constraints, with architectural variety explored subsequently. This extends decision-tree practice by attaching auditable artefacts to checkpoint outcomes.

4.3 Formative Industry Walk-through with First-time Users

As anticipated in Section 1, a formative walk-through was conducted with a commercial vendor of building-energy and HVAC design software, whose toolchain supports certification, compliance, and engineering workflows. The session followed the procedure outlined in Figure 1. Scenarios were presented sequentially. For each scenario, the facilitator read the relevant gate prompts, participants reasoned through the decision path, and the selected gate outcomes and evidence artefacts were recorded on a structured template.

The aim was to assess whether professionals encountering blockchain for the first time could use the framework to distinguish between cases that warrant further DLT analysis and cases that should exit towards simpler architectures. The resulting observations are descriptive and setting-bound; they are interpreted as evidence of practical intelligibility rather than as a controlled measurement of usability or effectiveness.

Participants (product management, application engineering, customer support) worked through two anonymised scenarios representative of the domain (Table 4). In the first scenario, a multi-party data exchange utilised a continuously online trusted intermediary and did not require cross-boundary immutability; an early, documented stop was reached at the necessity triage, and a conventional database with tamper-evident logging was selected. In the second, a constrained inter-organisational workflow required limited shared writes and verifiable audit trails; participants progressed beyond the triage, fixed the permission and disclosure choices, compiled the minimal evidence artefacts, and explicitly deferred platform selection until after implementability checkpoints.

Three observations are relevant to the practical adoption of the system. Firstly, the plain-language prompts and explicit off-ramps enabled first-time users to justify non-DLT outcomes without specialist jargon, aligning with the reject-first intent. Secondly, separating necessity from implementability prevented premature platform discussions: permission and consensus choices were postponed until legality, absorptive capacity, and SLO envelopes had been established. Thirdly, artefact templates for compliance and integration acted as scaffolding, rendering these considerations concrete from the outset and reducing ambiguity about operational readiness.

In phase-gate terms, Scenario 1 stopped at P3 (off-ramp: DB with tamper-evident logging); Scenario 2 progressed beyond P5–P6 with platform choice deferred to P7–P9.

4.4 Gate Ablation

This analysis is explicitly *non-empirical*: it transposes the ablation idiom from machine-learning evaluation to the checkpoint design, asking which failure modes each checkpoint precludes. No new data is produced; each removal is paired with existing evidence from the review corpus. This exercise establishes *argumentative necessity* relative to known failure modes, rather than causal contribution in the experimental sense. Removing one checkpoint at a time:

- **No compliance checkpoint (G0).** Decisions could proceed despite unlawful processing or irreconcilable records duties; public-sector evaluations identify such omissions as recurrent failure causes [6].
- **No necessity triage (G1–G4).** Early off-ramps would be absent where single-writer authority, trusted third parties, or a lack of cross-boundary immutability apply, contrary to canonical guidance that favours databases or tamper-evident logs in these settings [8, 29, 32].
- **No integration feasibility (G7).** Platform selection could be made in abstraction from identity, observability, runbooks, and change management (dimensions often scored but not gated in control-heavy frameworks [5, 13]).
- **No performance/finality under SLOs (G8).** Latency/finality budgets would be asserted rather than demonstrated against application envelopes, despite survey recommendations to avoid naïve extrapolation of vendor claims [5, 13].
- **No economics/incentives (G9).** Fee volatility, staffing, or misaligned incentives could surface after architectural commitment; comparative syntheses trace poor outcomes to underestimating steady-state costs and governance readiness [13, 29].

Each removal reveals a failure mode that has already been documented in the literature. Making the checkpoints non-compensatory means that the burden of proof is placed on minimal artefacts (e.g., a lawful-basis memo for compliance; an interface/runbook inventory for integration).

4.5 Threats to Validity and Limitations

Construct validity depends on whether the checkpoint definitions capture the intended distinction between blockchain necessity and blockchain implementability. This risk is mitigated by expressing each gate through explicit questions, reject conditions, and minimal evidence artefacts, although borderline cases involving partial trust, mixed governance, or hybrid architectures may still require expert interpretation.

Internal validity is limited by the formative nature of the walk-through and by the qualitative character of the ablation analysis. The walk-through was single-site and artefact-centred, while the ablation establishes argumentative necessity relative to known failure modes rather than causal effects. These limitations are appropriate for an early-stage design-science artefact, but they preclude claims of statistical effectiveness.

External validity is bounded by jurisdiction, sector, and organisational maturity. The legal and records-management assumptions are primarily framed around an EU context, and sector-specific obligations may alter the operation of G0 and G7 in other domains [6, 13]. Moreover, the review uses IEEE Xplore as its operational counting frame and English-language predicates, which may under-represent domain-specific or non-English adoption guidance.

Conclusion validity is shaped by the conservative design choice. Since the gates are non-compensatory, the framework may generate false negatives in borderline cases where a blockchain solution could become viable after substantial redesign. This bias is intentional and is judged defensible in regulated or public-interest settings, where unjustified blockchain adoption can impose costs, compliance exposure, and governance complexity before benefits are demonstrated [1].

5 Conclusion

This paper presents a conservative and auditable framework for deciding when *not* to use blockchain. The framework translates recurring adoption criteria into a compact phase-gate procedure comprising must-meet gates, minimal evidence artefacts, and documented off-ramps to non-DLT alternatives. By distinguishing necessity (lawfulness, independent writers under limited trust, and cross-boundary immutability) from implementability (permission and disclosure fit, integration readiness, performance and finality under SLOs, and economic sustainability), the framework prevents premature platform selection and makes rejection decisions explicit rather than implicit.

The analytical assessment shows that the framework preserves the reject-first logic of canonical decision trees while extending it with auditable compliance, integration, performance, and economic gates. The formative walk-through suggests that first-time users can follow the decision path, identify evidence artefacts, and justify early rejection when simpler architectures suffice. In social-good

Table 4: Formative walk-through: scenarios and outcomes.

Scenario	Salient conditions	Decisive answers	Outcome
#1 Multi-party data exchange	Trusted intermediary acceptable	TTP acceptable	Early stop at necessity triage; off-ramp: DB + tamper-evident logging.
#2 Constrained inter-organisational workflow	Limited shared writes; verifiable audit trail required	Independent writers present; immutability warranted; no continuously online TTP	Proceed beyond triage; permission/disclosure fixed; platform selection deferred to implementability checks.

and public-interest contexts, this is not merely a technical optimisation: avoiding unnecessary blockchain adoption protects scarce resources, reduces compliance exposure, and preserves institutional accountability.

Future work should validate the framework through multi-site studies, measure per-gate exit patterns, and develop tooling for automated evidence capture and permission-branching support. Further work should also examine hybrid designs that combine conventional databases with notarisation mechanisms, as well as bridge architectures that connect permissioned and permissionless infrastructures under explicit governance constraints.

Acknowledgments

This work was supported by the PRIN B4A project. The authors thank the participants involved in the formative walk-through for their time and feedback. The authors used ChatGPT (OpenAI) to improve the readability and clarity of selected sentences.

Data Availability

A Figshare companion bundle (DOI placeholder <https://doi.org/10.6084/m9.figshare.32112619>).

References

[1] T. Almeshal and A. Alhogail. 2021. Blockchain for Businesses: A Scoping Review of Suitability Evaluations Frameworks. 1 pages. doi:10.1109/ACCESS.2021.3128608

[2] Elli Androulaki, Artem Barger, Vita Bortnikov, Christian Cachin, Konstantinos Christidis, Angelo De Caro, David Enyeart, Christopher Ferris, Gennady Laventman, Yacov Manevich, Srinivasan Muralidharan, Chet Murthy, Binh Nguyen, Manish Sethi, Gari Singh, Keith Smith, Alessandro Sorniotti, Chrysoula Stathakopoulou, Marko Vukolić, Sharon Weed Cocco, and Jason Yellick. 2018. Hyperledger Fabric: A Distributed Operating System for Permissioned Blockchains. In *Proceedings of the Thirteenth EuroSys Conference (EuroSys '18)*. ACM, New York, NY, USA, 1–15. doi:10.1145/3190508.3190538

[3] I. Ar, Ismail Erol, Iskender Peker, A. Ozdemir, T. Medeni, and I. T. Medeni. 2020. Evaluating the feasibility of blockchain in logistics operations: A decision framework. 113543 pages. doi:10.1016/j.eswa.2020.113543

[4] Authors. 2026. Companion materials for "Blockchain or Not? An Auditable Decision Framework for Organisations": SLR protocol, phase/gate notation, flowcharts, and operational-tooling references. Companion to BANDIT '26 / GoodIT 2026.. doi:10.6084/m9.figshare.TODO_BANDIT_2026

[5] M. Belotti, Nikola Bozic, G. Pujolle, and Stefano Secci. 2019. A Vademecum on Blockchain Technologies: When, Which, and How. 3796-3838 pages. doi:10.1109/COMST.2019.2928178

[6] Diego Cagigas, Judith Clifton, Daniel Díaz-Fuentes, Marcos Fernández-Gutiérrez, and C. Harpes. 2023. "Blockchain in government: toward an evaluation framework". 397 - 414 pages. doi:10.1080/25741292.2023.2230702

[7] Vittorio Capocasale, Maria Elena Bruni, and Guido Perboli. 2025. Technological insights on blockchain adoption: the electric vehicle supply chain use case. doi:10.1108/ejim-06-2024-0701

[8] M. Chowdhury, A. Colman, M. A. Kabir, Jun Han, and Paul Sarda. 2018. Blockchain Versus Database: A Critical Analysis. 1348-1353 pages. doi:10.1109/TrustCom/BigDataSE.2018.00186

[9] Robert G. Cooper. 2008. Perspective: The Stage-Gate® Idea-to-Launch Process—Update, What's New, and NexGen Systems. *Journal of Product Innovation Management* 25, 3 (2008), 213–232.

arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-5885.2008.00296.x> doi:10.1111/j.1540-5885.2008.00296.x

[10] P. H. B. Correia, Marco A. Marques, Marcos A. Simplicio, Luciano Ermlivitch, Charles C. Miers, and M. A. Pillon. 2024. Comparative Analysis of Permissioned Blockchains: Cosmos, Hyperledger Fabric, Quorum, and XRPL. 464-469 pages. doi:10.1109/Blockchain62396.2024.00068

[11] Thao T. A. Dinh, Ji Wang, Gang Chen, Rui Liu, Beng Chin Ooi, and Kian-Lee Tan. 2017. BLOCKBENCH: A Framework for Analyzing Private Blockchains. *Proceedings of the ACM International Conference on Management of Data (SIGMOD) – arXiv version* (2017). <https://arxiv.org/abs/1703.04057>

[12] Siamak Farshidi, S. Jansen, Sergio España, and Jacco Verkleij. 2020. Decision Support for Blockchain Platform Selection: Three Industry Case Studies. 1109-1128 pages. doi:10.1109/TEM.2019.2956897

[13] S. Gourisetti, M. Mylrea, and H. Patangia. 2020. Evaluation and Demonstration of Blockchain Applicability Framework. 1142-1156 pages. doi:10.1109/TEM.2019.2928280

[14] Florian Gräbe, Niclas Kannengießer, S. Lins, and A. Sunyaev. 2020. Do Not Be Fooled: Toward a Holistic Comparison of Distributed Ledger Technology Designs. 10 pages. doi:10.24251/hicss.2020.770

[15] Vikas Hassija, S. Zeadally, Ishan Jain, Aman Tahiliani, V. Chamola, and Shashank Gupta. 2021. Framework for determining the suitability of blockchain: Criteria and issues to consider. doi:10.1002/ett.4334

[16] Marko Hribernik, Kathrin Zero, S. Kummer, and D. Herold. 2020. City logistics: Towards a blockchain decision framework for collaborative parcel deliveries in micro-hubs. doi:10.1016/j.trip.2020.100274

[17] Jens J. Hunhevicz and D. Hall. 2020. Do you need a blockchain in construction? Use case categories and decision framework for DLT design options. 101094 pages. doi:10.1016/j.aei.2020.101094

[18] Lahlou Imane, Motaki Noureddine, Sarsri Driss, and L'yarfi Hanane. 2023. Towards Blockchain-Integrated Enterprise Resource Planning: A Pre-Implementation Guide. 11 pages. doi:10.3390/computers13010011

[19] T. Koens and E. Poll. 2018. What Blockchain Alternative Do You Need? 113-129 pages. doi:10.1007/978-3-030-00305-0_9

[20] Ben Laurie, Adam Langley, and Emilia Kasper. 2013. Certificate Transparency. RFC 6962 (Proposed Standard). <https://www.rfc-editor.org/rfc/rfc6962>

[21] Ben Laurie, Eran Messeri, and Rob Stradling. 2021. Certificate Transparency Version 2.0. RFC 9162 (Standards Track). <https://www.rfc-editor.org/rfc/rfc9162>

[22] Kun Li, Ashish Rajendra Sai, and V. Urovi. 2024. Do you need a blockchain in healthcare data sharing? A tertiary review. doi:10.37349/edht.2024.00014

[23] Sin Kuang Lo, Xiwei Xu, Yin Kia Chiam, and Q. Lu. 2017. Evaluating Suitability of Applying Blockchain. 158-161 pages. doi:10.1109/ICECCS.2017.26

[24] Jan Mendling, Ingo Weber, Wil Van Der Aalst, Jan Vom Brocke, Cristina Cabanillas, Florian Daniel, Søren Debois, Claudio Di Ciccio, Marlon Dumas, Schahram Dustdar, Avigdor Gal, Luciano García-Bañuelos, Guido Governatori, Richard Hull, Marcello La Rosa, Henrik Leopold, Frank Leymann, Jan Recker, Manfred Reichert, Hajo A. Reijers, Stefanie Rinderle-Ma, Andreas Solti, Michael Rosemann, Stefan Schulte, Munindar P. Singh, Tijs Slaats, Mark Staples, Barbara Weber, Matthias Weidlich, Mathias Weske, Xiwei Xu, and Liming Zhu. 2018. Blockchains for Business Process Management - Challenges and Opportunities. *ACM Trans. Manage. Inf. Syst.* 9, 1, Article 4 (Feb. 2018), 16 pages. doi:10.1145/3183367

[25] Ahmed Afif Monrat, O. Schelén, and K. Andersson. 2022. Applicability Analysis of Blockchain Technology. 8 pages. doi:10.1109/CSDE56538.2022.10089307

[26] Stefan Naef, Stephan M. Wagner, and C. Saur. 2022. Blockchain and network governance: learning from applications in the supply chain sector. 932 - 946 pages. doi:10.1080/09537287.2022.2044072

[27] Matthew J. Page, Joanne E. McKenzie, Patrick M. Bossuyt, Isabelle Boutron, Tammy C. Hoffmann, Cynthia D. Mulrow, Larissa Shamseer, Jennifer M. Tetzlaff, Elie A. Akl, Sue E. Brennan, Roger Chou, Julie Glanville, Jeremy M. Grimshaw, Asbjørn Hróbjartsson, Manoj M. Lalu, Tianjing Li, Elizabeth W. Loder, Evan Mayo-Wilson, Steve McDonald, Luke A. McGuinness, Lesley A. Stewart, James Thomas, Andrea C. Tricco, Vivian A. Welch, Penny Whiting, and David Moher. 2021. The PRISMA 2020 statement: an updated guideline for reporting systematic reviews. *BMJ* 372 (2021), n71. doi:10.1136/bmj.n71

[28] M. Peck. 2017. Blockchain world - Do you need a blockchain? This chart will tell you if the technology can solve your problem. 38-60 pages. doi:10.1109/MSPEC.

- 2017.8048838
- [29] Brian A. Scriber. 2018. A Framework for Determining Blockchain Applicability. 70-77 pages. doi:10.1109/MS.2018.2801552
- [30] A. Simaremare, I. Aditya, F. Haryadi, and H. Indrawan. 2021. Suitability study of Blockchain application in electric utility company business processes. doi:10.1088/1757-899X/1098/5/052105
- [31] M. Touloupou, Klitos Christodoulou, Antonios Inglezakis, Elias Iosif, and Marinos Themistocleous. 2021. Towards a Framework for Understanding the Performance of Blockchains. 47-48 pages. doi:10.1109/brains52497.2021.9569810
- [32] K. Wüst and Arthur Gervais. 2018. Do you Need a Blockchain? 45-54 pages. doi:10.1109/CVCBT.2018.00011
- [33] Xiwei Xu, I. Weber, M. Staples, Liming Zhu, J. Bosch, L. Bass, Cesare Pautasso, and Paul Rimba. 2017. A Taxonomy of Blockchain-Based Systems for Architecture Design. 243-252 pages. doi:10.1109/ICSA.2017.33