Information Systems 140 (2026) 102723

Contents lists available at ScienceDirect

Information Systems

journal homepage: www.elsevier.com/locate/is

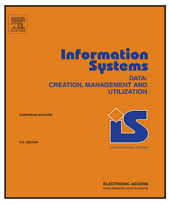

Designing a blockchain-based information system for procurement processes—Balancing decentralization, scalability, and security while maintaining privacy☆

Valeriya Arnold [a,b,c,d], Tobias Guggenberger [a,b,d], Jan Stramm [a,b,c,d,*], Nils Urbach [a,c,d]

[a] Fraunhofer Institute for Applied Information Technology FIT, Branch Business and Information Systems Engineering, Germany
[b] University of Bayreuth, Faculty III Law, Business and Economics, Germany
[c] Frankfurt University of Applied Sciences, Faculty 3 Business and Law, Germany
[d] FIM Research Center for Information Management, Germany

ARTICLE INFO

Keywords:
Blockchain
Procurement
Blockchain Trilemma
Privacy
Design Science Research
Zero Knowledge Proofs

ABSTRACT

The inter-organizational processes in procurement remain burdened by media discontinuity, inefficiencies, and a lack of trust among trading partners. Blockchain-based information systems are frequently proposed as a remedy because they enable shared, tamper-evident records. However, existing instantiations rarely scale beyond small consortia because they fail to address the extended blockchain trilemma, which requires simultaneously achieving decentralization, scalability, security, and strict privacy requirements for sensitive commercial information. In contrast to prior blockchain procurement prototypes that manage the extended blockchain trilemma primarily through permissioned architectures, this study investigates how a blockchain-based information system can be designed to reconcile the trade-offs inherent in the extended trilemma, achieving a viable balance through architectural allocation and cryptographic enforcement. Following the design-science research paradigm, an empirically validated problem statement is synthesized from a structured literature review and expert interviews. Five design objectives are derived, evaluated, and used to guide a prototype design, which is then iteratively refined and evaluated through quantitative and formative assessments and sixteen semi-structured expert interviews. Reflection on the build–evaluate cycles yields two design principles: (1) Balancing decentralization, scalability, and security by using a public chain as a trust anchor, a Layer 2 for scaling, and decentralized communication between layers. (2) Maintaining that balance when privacy is required by integrating efficient, resilient cryptography and minimizing control points. These principles extend existing procurement research, linking business requirements to infrastructural choices, providing a transferable foundation for scholars and practitioners aiming to deploy secure, scalable, and privacy-preserving blockchain solutions in inter-organizational contexts.

1. Introduction

Procurement processes serve as critical information exchange mechanisms where organizations convert internal demand signals into contracts, orders, and payments, yet recent disruptions from the COVID-19 pandemic or cyber-attacks on firms such as Maersk have revealed fundamental weaknesses in the underlying technological infrastructure supporting these operations [1–4]. As one of the most information-intensive business processes, procurement accounts for substantial operating costs and requires continuous inter-organizational exchange of highly sensitive commercial data, including unit prices, contractual terms, and negotiation records [5–8]. These processes, however, remain fragmented across incompatible information systems, vulnerable to manual errors, and exposed to opportunistic behavior from trading partners. The resulting lack of transparency and erosion of mutual trust between organizations inflates transaction cycle times, generates frequent disputes, and systematically undermines operational resilience [9–12].

Researchers and practitioners have increasingly developed blockchain technology as a technical remedy for these shortcomings, recognizing its potential to address fundamental trust and transparency challenges. It offers a replicated, tamper-evident ledger that can be shared among mutually distrusting stakeholders, such as legally independent firms involved in procurement [13–15]. A shared ledger

☆ This article is part of a Special issue entitled: 'BISs' published in Information Systems.

* Corresponding author.
E-mail address: jan.stramm@fit.fraunhofer.de (J. Stramm).

https://doi.org/10.1016/j.is.2026.102723
Received 29 September 2025; Received in revised form 9 March 2026; Accepted 24 March 2026
Available online 27 March 2026

could remove many reconciliation loops that dominate today's procurement processes. Smart contracts might automate routine steps such as delivery confirmation and payment release, thereby shortening cycle times [16]. The same immutable record is expected to simplify audit and reporting duties, and tokenized settlement mechanisms promise near-real-time transfer of value without an additional intermediary layer [17]. Yet large-scale adoption has lagged far behind initial expectations. High-profile pilots such as TradeLens and Everledger moved beyond the proof-of-concept stage but were ultimately wound down after failing to secure sufficient stakeholder participation and after disputes over data ownership and governance [18,19]. Other efforts, including IBM and Walmart's Food Trust as well as De Beers' Tracr, continue to operate but remain limited to tightly scoped use cases [20,21]. Although these projects yield valuable design knowledge on how blockchain can serve procurement-related information flows, their restricted uptake has led to questioning the technology's broader relevance, noting that meaningful traction persists mainly in payment and other token-centric financial applications [16,22].

As outlined by Mishunin (2024), one of the possible reasons for limited adoption and successful outcomes may stem from fundamental design tensions inherent in blockchain architectures [23]. The canonical *"trilemma"* holds that decentralization, scalability, and security cannot be maximized simultaneously [24,25]. When blockchain technology is applied in enterprise or industrial contexts, privacy becomes a fourth design dimension alongside decentralization, scalability, and security, yielding what we refer to as the extended blockchain trilemma. The demand for confidentiality is particularly acute in procurement, where unit prices and contract terms are highly sensitive and form the basis of competitive advantage. Under these conditions, the full transparency and limited throughput of public permissionless networks such as Ethereum are untenable, as firms must shield commercial details and potentially process thousands of transactions per hour. As a consequence, most industrial pilots have turned to permissioned platforms such as Hyperledger Fabric, Quorum, or Corda to protect sensitive data [26–28]. While these systems improve privacy and performance, they re-create central control points, complicate multi-party governance, and dilute the open verifiability that initially motivated interest in blockchain [23,29]. Current architectural approaches thus continue to exhibit fundamental tensions between competing design requirements, constraining blockchain deployments in procurement to niche consortia or narrowly defined use cases [28,30]. What is still missing is a design that combines the decentralized trust, security, and auditability of public permissionless blockchains, serving as a trust anchor in the sense of a publicly verifiable and tamper-resistant base layer, with the scalability and privacy required for routine, data-sensitive procurement operations, found in private permissioned settings. Prior blockchain procurement studies provide valuable insights into processes and prototypes, yet most artifacts address only subsets of these tensions. In particular, privacy and scalability are frequently achieved by restricting participation through permissioned or consortium architectures, which can be effective in narrowly scoped networks but tend to reintroduce governance dependencies and control points that limit open verifiability beyond the consortium boundary. Conversely, public permissionless deployments preserve decentralization and auditability but typically struggle with the confidentiality and transaction volumes required for routine procure-to-pay execution. As a result, existing work offers limited guidance on how to design procurement information systems that preserve permissionless trust and auditability while enabling private and scalable execution of sensitive procurement workflows.

Emerging cryptographic innovations are beginning to address these fundamental design tensions in blockchain architectures for enterprise applications. Cryptographic primitives such as Zero-Knowledge Proofs (ZKPs) [31], homomorphic encryption [32], and roll-up architectures [33] promise confidentiality and throughput while retaining the decentralization and security of public networks. Early empirical evidence supports this view: The TEN-Network employs a trusted execution environment (i.e., a hardware-isolated environment that protects code execution and data from external interference) paired with an optimistic roll-up architecture to secure sensitive data [34], while Aztec applies ZKPs to enable private, scalable transfers on Ethereum [35]. These developments challenge the prevailing assumption that only private permissioned blockchains can meet business requirements regarding privacy and scalability for procurement. However, guidance on integrating such primitives into a viable blockchain design for routine procurement workflows is sparse. We thus see a gap in the literature as well as design knowledge that the extended blockchain trilemma needs to be reconciled, i.e., achieving a viable balance through design choices that allocate responsibilities to the layer where they can be satisfied, rather than claiming to eliminate the underlying trade-offs. This gap motivates our research question:

RQ: How to design a blockchain-based information system for procurement processes that balances decentralization, scalability, and security while maintaining privacy?

To investigate this question, we followed the Design-Science Research (DSR) paradigm [36]. This approach is particularly suitable because addressing the extended blockchain trilemma in procurement processes requires the purposeful construction and evaluation of a novel system artifact under real-world technical, regulatory, and organizational constraints. First, a structured literature review yielded a concise problem statement that combined procurement business needs with blockchain design tensions, along with a preliminary set of Design Objectives (DOs). These were then reviewed, refined, and evaluated by experts. Guided by these DOs, we constructed a prototype that underwent iterative presentation [37], formative testing, and a summative evaluation involving quantitative benchmarking and sixteen expert interviews [38]. The experts' feedback was coded against the DOs, allowing us to assess and further adjust the design. Reflection on the build–evaluate cycles yielded two prescriptive Design Principles (DPs) that extend current knowledge beyond the specific artifact [39]. This study makes four contributions. First, it articulates an empirically grounded problem statement that links procurement's information-exchange requirements to the extended blockchain trilemma. Second, it derives and evaluates five design objectives that translate these tensions into implementable requirements for blockchain-based procurement information systems. Third, it designs and demonstrates a layered prototype that anchors trust (i.e., globally verifiable integrity) and auditability on a public permissionless L1 (i.e., the base blockchain providing shared consensus and immutability) while executing confidential, high-volume procurement workflows on a L2 roll-up using ZKPs (i.e., a scaling layer that processes transactions off-chain while inheriting Layer 1 security guarantees). Fourth, it distills two prescriptive design principles that connect procurement requirements to infrastructural design choices, informing blockchain-based inter-organizational information systems beyond the specific prototype. These principles provide actionable guidance for researchers and practitioners designing blockchain solutions for inter-organizational information sharing in procurement and related applications.

2. Background

2.1. Information exchange in procurement processes

Procurement transforms an internal demand signal into a delivered good and a settled payment. Regardless of whether the workflow is labeled "procure-to-pay", "source-to-settle", or "purchase-to-invoice", it typically unfolds in five phases: need identification, supplier selection, ordering, fulfillment, and settlement [1,5–8]. In the remainder of this paper, we will use the wording "procure-to-pay" to refer to this set of phases and underlying processes. Each phase generates data that must cross organizational boundaries, yet the corresponding information exchange is often mediated by isolated software tools and

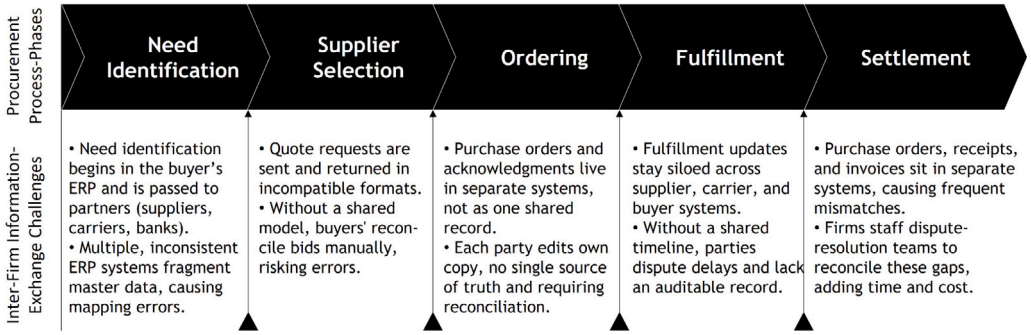

Fig. 1. Procurement process-phases and information exchange challenges.

heterogeneous platforms, which undermines overall efficiency [9–12]. The following section examines each phase in detail and identifies the key shortcomings that arise during inter-firm information exchange (Fig. 1).

First, *need identification* begins inside the buyer's Enterprise-Resource-Planning (ERP) environment. Yet external parties later consume this data as well (e.g., material numbers, cost centers, delivery dates). Since large firms often run multiple Enterprise-Resource-Planning (ERP) instances with diverging master-data conventions, the very first dataset is already fragmented [40]. Inconsistent item codes or unit measures resurface as data exceptions once suppliers, carriers, or banks try to map the request to their own catalogs.

During *supplier selection*, the requisition is recast as a request for quotation or proposal. The document crosses the corporate firewall and returns in as many formats as there are bidders, e.g., PDFs, spreadsheets, and portal exports. Lacking common semantics, buyers often manually translate and standardize the offers according to their internal needs, which obscures provenance and introduces transcription risk [41]. The absence of a shared data model for reference, rather than a particular messaging tool, can be the root of friction here.

Ordering formalizes the commercial commitment through a purchase order and an order acknowledgment [8]. Conceptually, these should be two views of one transaction dataset. In practice, they are separate records stored in isolated systems. No canonical version exists if either party modifies quantity, price, or delivery date. The problem is not how the documents travel via email, portal, or Application Programming Interface (API), but that each side maintains its own mutable ledger. Hence, reconciliation becomes a permanent overhead [9].

The subsequent *fulfillment* phase involves status updates from manufacturers, carriers, and warehouses. Even when updates occur in near-real time, they remain siloed. A "goods shipped" event in a supplier portal does not automatically become a "goods in transit" event in the buyer's system [5,6]. Without a shared, time-stamped timeline, supply chain stakeholders dispute whether delays are factual or merely a visibility gap, and they often lack an auditable record on which to base decisions.

Settlement should close the loop through a three-way match between the purchase order, goods receipt, and invoice [7]. Yet these three artifacts reside in separate databases maintained by two organizations, so any earlier divergence reappears as a mismatch that must be resolved case by case [12]. Over time, the volume and complexity of these exceptions have led many firms to create dedicated "dispute-resolution" or "claims-management" teams that mediate between procurement, logistics, and accounts payable, making it an unofficial sixth phase that absorbs additional cycle time and cost [42,43]. The labor devoted to such reconciliations reflects not a payment-technology shortfall but merely the absence of an immutable, mutually trusted transaction history.

The five procurement phases as shown in Fig. 1 expose a standard set of structural impediments regarding the information exchange in procurement processes: fragmented system landscapes, non-standard data semantics, and isolated ledgers that seldom agree. The preceding section has illustrated the issues primarily through the data flow between a single buyer and a single supplier. In reality, most supply chains involve multiple tiers of subcontractors, logistics providers, and financial institutions. Each additional party adds further ERP instances, data models, and governance rules, which multiply semantic gaps and points of failure [12]. Even sophisticated ERP suites and their satellite portals accelerate internal workflows only up to the point where data must cross the organizational boundary [11,40]. Beyond that boundary, manual translation, duplicate storage, and after-the-fact reconciliation remain common practice. The resulting friction is more than an operational nuisance: with procurement accounting for significant corporate spend, every delay or mismatch inflates working-capital locks, invites compliance risk, and weakens inter-firm trust [1,8].

2.2. Blockchain-based information exchange in procurement

Distributed Ledger Technology (DLT), particularly blockchain, has been advanced as a structural remedy for the coordination deficits that characterize inter-firm procurement [13]. A blockchain is a replicated, append-only database whose entries are collectively validated by a network of nodes rather than a single platform operator [44]. Each block contains a cryptographic reference to its predecessor, yielding an immutable event history. At the same time, a consensus algorithm (e.g., Proof-of-Stake, Byzantine fault-tolerant voting) ensures that all honest nodes arrive at the same state even in the absence of mutual trust [45,46]. Smart contracts, deterministic programs stored on the ledger, extend this functionality by executing agreed business logic once predefined conditions are met [47].

For enterprise use, several architectural choices exist along a continuum from public, permissionless networks (e.g., Ethereum Mainnet) to strictly permissioned, consortium-governed ledgers (e.g., Hyperledger Fabric or Corda) [48]. Permissionless blockchains allow open participation without centralized admission control, whereas permissioned blockchains restrict participation to a predefined set of approved entities. Hybrid designs, side-chains, and L2 roll-ups offer additional levers for balancing openness, throughput, and confidentiality [33]. Regardless of the variant selected, three generic blockchain capabilities are especially relevant for procurement: shared state, time-stamped immutability, and programmable execution [2,22,49]. A *shared state* guarantees that all parties reference the same canonical ledger entry, be it a purchase order, shipment notice, or invoice, thereby eliminating potential multiple versions of data that proliferate across isolated ERP databases. Furthermore, *time-stamped immutability* can ensure that every ledger update is irrevocably linked to its cryptographic past, thus allowing for the dissolution of later disputes about who changed what and when by inspection rather than negotiation [28]. *Programmable execution* provided by smart contracts can additionally automate conditional actions such as release of advance payment after an on-chain delivery confirmation, automatic late-fee calculation, or self-assessment

of value-added tax, without recourse to manual three-way matching or third-party escrow services [50].

These properties map directly onto the structural impediments outlined in the previous sub-section. Fragmented system landscapes become federated around a single ledger: Non-standard semantics can be replaced by a jointly maintained data schema that every transaction must satisfy before it is committed. The need for ex-post reconciliation gives way to real-time co-validation, fostering trust without a central clearing house. Early pilots report shorter cycle times, fewer invoice disputes, and improved audit readiness, especially in multi-tier supply chains where intermediate actors had previously remained opaque [26,43,51,52]. Use cases within procurement processes include digital supplier credentials that speed up onboarding processes, on-chain contract management [20], provenance verification for high-value goods [21], conditional payments triggered by verified milestones [53], and automated compliance reporting. In most designs, the ledger supplements rather than replaces existing ERP suites, serving as a synchronization layer that records cryptographic proofs or limited data extracts while keeping sensitive operational details within corporate boundaries.

2.3. Trilemma and privacy design challenges

Distributed ledgers offer tamper-evident data sharing, yet they must be designed within a space shaped by mutually constraining objectives [54]. This circumstance coined *blockchain trilemma* states that decentralization, scalability, and security cannot be maximized simultaneously: Broad participation and strong fault tolerance raise consensus overhead, whereas high throughput usually demands fewer validators or weaker consistency guarantees [45]. Procurement intensifies the dilemma by adding a fourth axis, commercial privacy, because unit prices and contract terms are highly sensitive and underpin competitive advantage [25,30]. Therefore, any blockchain architecture for procurement processes must balance four goals rather than three.

Public permissionless networks like Ethereum excel at decentralization and data integrity, but achieve only 15–30 transactions per second, far below the tens of thousands of daily events typical in the global procurement space [55]. Their complete transparency is also unacceptable to firms that must keep commercial details confidential. To address throughput and confidentiality, many pilots adopt permissioned or consortium blockchains (e.g., Hyperledger Fabric, Quorum, Corda). Restricting the validator set this increases performance and enables access control for sensitive data. Still, it also reintroduces central points of failure, shifts trust to governance bodies, and fractures the openness that initially motivated blockchain adoption [29]. Therefore, effective governance covering validator admission, protocol upgrades, and dispute resolution is critical. Without it, networks tend to fragment into non-interoperable silos [48]. A range of technical mitigations has emerged. Notary-node schemes (as in Corda) prevent double spending with minimal messaging but concentrate trust in a single entity [56]. Byzantine-fault-tolerant protocols distribute trust yet incur quadratic communication costs as validator numbers rise. Sharding, side-chains, and L2 roll-ups move transaction processing off-chain and settle aggregates on a secure base layer, but each technique imports new trust assumptions, such as honest shard majorities or reliable fraud proofs, that are hard to guarantee across independent buyers and suppliers [33].

Privacy remains the most challenging axis. In procurement, indiscriminate transparency can reveal price lists, bill-of-materials data, or counterparty relationships [54,57]. Public-key pseudonymity offers limited protection, and permissioned ledgers allow validators to inspect clear-text payloads [58]. Researchers therefore explore privacy-enhancing technologies, such as ZKPs, which enable private computation by allowing one party to prove the validity of a statement without revealing the underlying data [31], homomorphic commitments [32], or secure multi-party computation [59]. From an information systems perspective, these cryptographic mechanisms are not ends in themselves but serve as design enablers, allowing organizations to verify process compliance and transaction validity without disclosing sensitive operational or commercial data to other network participants. However, they also add computational overhead, increasing integration complexity and introducing new attack surfaces. Finally, this finding adds to trilemma tensions and helps explain why blockchain pilots in procurement seldom progress beyond proof of concept. Architectures that scale often centralize validation, while designs that ensure confidentiality rely on cryptography, which hampers performance. Reconciling decentralization, scalability, security, and privacy within one system remains the central challenge that motivates our artifact and the design principles derived from it.

3. Method

This study follows the six-step DSR process: Problem Identification and Motivation; Definition of Design Objectives; Design, Development, and Demonstration; Evaluation and Communication, as proposed by Peffers et al. (2007) [36]. This paradigm is particularly suitable because the research question focuses on *how* to design an information system artifact that reconciles competing infrastructural requirements, rather than on explaining adoption behavior or evaluating an existing system ex post. Alternative empirical approaches, such as surveys or econometric analyses, cannot fully address this design-oriented research endeavor, as no comparable artifact currently exists in practice. Likewise, a purely technical or engineering approach would risk neglecting the organizational, regulatory, and inter-organizational constraints that shape procurement processes. The DSR framework links rigorous empirical analysis with the iterative construction and evaluation of artifacts, enabling us to ground design decisions in both scholarly knowledge and practitioner insights while ensuring academic validity and practical relevance. Regarding the communication of the research findings, we understand that this paper disseminates the results and extends scholarly knowledge, particularly in the context of the extended blockchain trilemma. See Fig. 2 for an overview of principal activities, methods, and artifacts.

Initially, to identify and motivate the problem, we carried out a SLR that followed the procedural guidance of Webster and Watson (2002) [60]. The review aimed to consolidate and organize the fragmented literature on blockchain-enabled procurement, establishing a clear problem understanding to guide the development of a conceptual architecture and prototype. Specifically, it produced a problem statement and an initial set of DOs for our new blockchain-based artifact [36]. We restricted the search to academic, peer-reviewed outlets. Using the search string *(“Procurement” OR “Procure-to-Pay” OR “Order-to-Cash”) AND (“Blockchain” OR “DLT” OR “Smart Contracts”) AND (“Prototype” OR “Artefact” OR “Proof of Concept” OR “Case Study”)* we queried the databases *IEEE Xplore, Science Direct, AIS eLibrary, ACM, Web of Science, SpringerLink and EconBIZ* on *[12.11.2024]*. This query returned $N_{4,218}$ publications.

We defined three inclusion criteria: (1) studies that report design knowledge on blockchain applications in procurement, (2) studies that address design knowledge for procure-to-pay operations (need identification, supplier selection, ordering, fulfillment, and settlement), and (3) empirical work on the integration of blockchain with procurement systems such as ERP. We excluded items that (a) offered only abstract claims about blockchain’s generic potential without technical detail, (b) provided general arguments for blockchain adoption in supply-chain management, (c) focused on supply-chain domains other than procurement (for example, distribution, logistics, manufacturing, or retailing), (d) were not written in English or German, or (e) were unavailable in full text. Backward and forward snowballing did not yield additional items that satisfied these criteria. Title-and-abstract screening followed by full-text review reduced the sample to a final set of N_{34} papers, 15 theory-oriented and 19 practice-oriented procurement design papers.

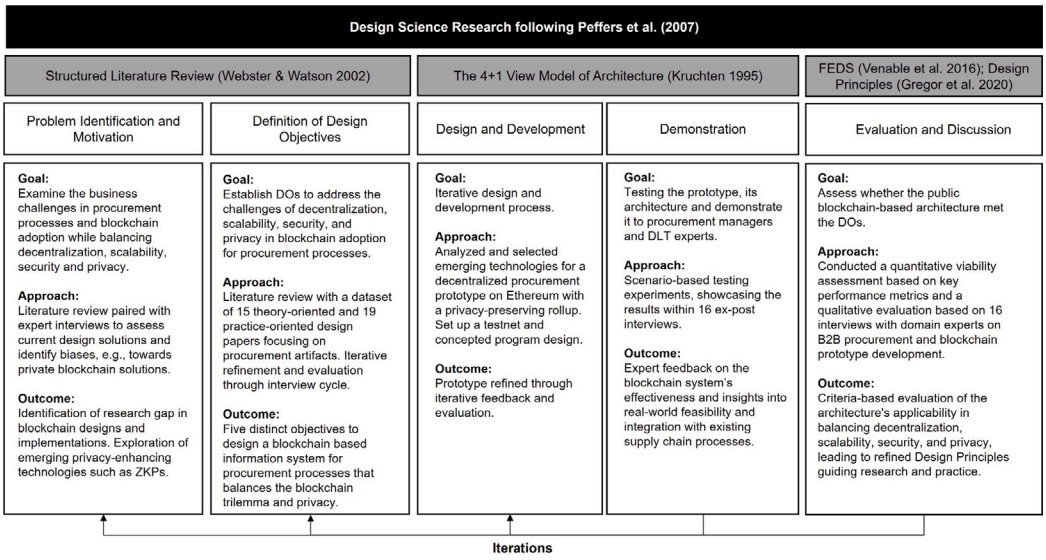

Fig. 2. Methodical approach based on Peffers et al. (2007) [36].

The sample was then coded using the qualitative analytics function of Citavi 7. The complete procedure is illustrated in Fig. 3.

Because academic literature alone could not capture emerging industrial practice, we complemented the SLR with insights from 16 semi-structured interviews that were iteratively integrated with the literature findings. Both sources of evidence revealed that current initiatives overwhelmingly rely on permissioned blockchains, aiming to solve privacy and limit the number of participants to avoid scalability issues, thus circumventing the blockchain trilemma, which is further complicated by privacy in general. These observations led to the problem statement that motivates our research: The need for a permissionless, privacy-preserving procurement platform capable of balancing the trilemma without compromising compliance or operational efficiency. Guided by this statement, we next defined the DOs of the new artifact. We therefore referred to our evidence from the SLR which resulted in the first set of derived DOs which were later confirmed and expanded by evidences from expert interviews taking into account the specificity of business needs as well as real-world implementation knowledge. As a result, we established five evaluated DOs.

Our artifact's design and development started with identifying technological configurations capable of meeting the blockchain trilemma and privacy requirements, thereby establishing a sound foundation for all subsequent Design Objectives (DOs) building upon it. We systematically examined public, permissionless blockchain platforms (Ethereum, Solana, Polkadot, and Findora) and complementary scaling and privacy-preserving technologies (including Fhenix, Calyx, Zkopru, and Aztec) with regard to their ability to simultaneously provide strong privacy guarantees and alleviate the fundamental constraints of blockchain systems. Most of these candidates, however, address only subsets of these requirements or introduce architectural trade-offs that conflict with our design goals. In particular, approaches based on optimistic verification entail delayed finality and temporary data exposure during challenge periods, while other platforms either lack sufficiently strong privacy guarantees or rely on comparatively immature and complex cryptographic assumptions [61]. As a result, Aztec emerged as the most suitable foundation for our artifact at the time of design, as it was the only solution in our analysis that combined privacy-preserving smart contracts with a rollup-based architecture explicitly designed to address the fundamental scalability, security, and decentralization constraints of public blockchains. We therefore focused on the Aztec privacy-first rollup protocol built on the Ethereum blockchain [62]. This choice allowed us to assess how much a state-of-the-art privacy-preserving, permissionless blockchain can satisfy the outlined requirements. Using Kruchten's 4 + 1 architectural view model, we produced logical and process views of the solution [37] and implemented a prototype in three development iterations. Each cycle addressed functionality gaps that surfaced during expert walkthrough, resulting in a procurement artifact that tokenizes purchase orders, batches transactions into rollups, and automates milestone-based payments via smart contracts. For demonstration purposes, we used the developed views to show a conceptual view and test results from the Docker-based environment, where one client ran Ethereum while another client implemented an L2 solution, in scenario-based experiments that replicated typical procurement workflows such as order issuance, goods receipt confirmation, and conditional payment release. Representatives (see Table 1) from the domains of procurement and supply chain management, finance, blockchain engineering, and information systems research observed these demonstrations and confirmed that the system operated as intended and aligned with prevailing business practices.

To evaluate the developed artifact, we followed the FEDs framework proposed by Venable et al. (2016), combining formative and summative evaluation activities across the design and demonstration cycles [38]. In line with DSR principles, the evaluation does not aim at statistical generalization, but at assessing whether the artifact plausibly satisfies its design objectives and can operate under realistic organizational and technical constraints [36]. During the development phase, formative evaluation focused on technical verification and explanatory analysis. This included validating core system functions such as token minting on L1, transfer and execution on L2, encrypted negotiation and transaction processing, and settlement back to L1 (cf. Design and Development section). In addition, error testing was conducted to identify and resolve issues related to transaction flow, data consistency, and cryptographic operations. The explanatory analysis further examined whether the use of ZKPs effectively protected sensitive procurement data while preserving auditability, traced the data path across all procure-to-pay steps, and confirmed that no business-critical information was revealed beyond what the process required.

The post-development evaluation consists of two complementary components, mirrored in the structure of the Evaluation section: a quantitative viability evaluation followed by a qualitative expert-based evaluation. The quantitative evaluation assesses whether the artifact can, in principle, sustain realistic procurement workloads in terms of throughput, latency, and economic feasibility. Given that the prototype operates in a localized Docker-based environment and that the underlying Layer-2 infrastructure is still evolving, this assessment is

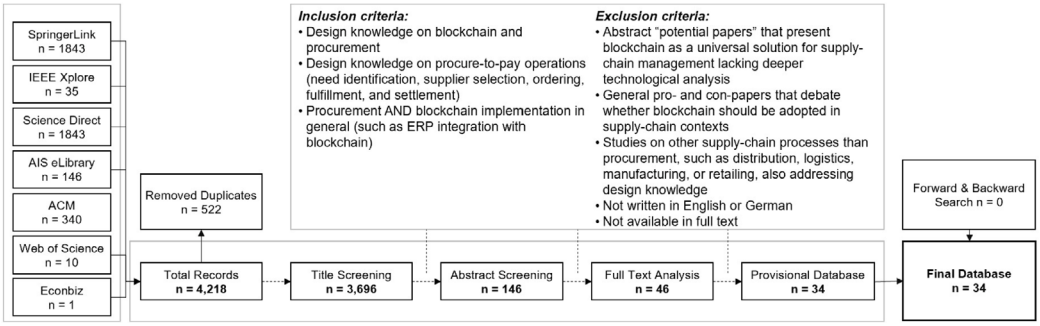

Fig. 3. Overview of the SLR based on Webster and Watson (2002) [60].

explicitly framed as a viability analysis rather than as definitive performance benchmarking. Controlled sandbox experiments, combined with workload projections and publicly available protocol characteristics, are used to derive indicative scalability targets and to identify potential structural bottlenecks. The qualitative evaluation provides a summative assessment of the artifact's conceptual soundness, practical relevance, and alignment with the predefined design objectives. It is based on 16 semi-structured expert interviews with practitioners and researchers from procurement, blockchain engineering, cryptography, and information systems (see Table 1). Each interview was conducted via a face-to-face video call that included a live demonstration of the artifact and lasted between 58 and 125 min (mean 81), yielding a total of 1299 min of recorded material. Recordings were transcribed using Microsoft Teams' speech-to-text service and subsequently manually verified and edited, resulting in 834 pages of textual data. Interview data were analyzed using a structured, theory-informed qualitative coding approach. Coding was guided by the predefined design objectives derived from the literature and earlier interview phases, while allowing for refinement based on empirical observations. Initial descriptive codes captured concrete expert assessments of the artifact (e.g., transparency gains, coordination overhead, latency sensitivity). These codes were then consolidated into higher-level analytical categories corresponding to the DOs, enabling a criteria-based evaluation of how well the artifact addresses each objective. While the coding process draws on grounded theory techniques (i.e., [63]) such as open and axial coding, it does not aim to generate inductive theory but to systematically organize expert judgments in relation to predefined evaluative criteria. The resulting evaluations are summarized in Table 5 and discussed in detail in Section 7.

In the final reflection and formalization step of Peffers et al.'s design-science process [36], we synthesized the empirical evidence gathered during the designing, developing, testing, and evaluating phases. The data confirmed that the artifact balances decentralization, scalability, and security while preserving the privacy and auditability required in sensitive procurement processes. We then drew on Gregor et al. (2020) to abstract these findings from the build-evaluate cycles beyond the concrete prototype and articulated two prescriptive design principles for employing permissionless, privacy-preserving blockchains [39]. These principles expand the academic knowledge base and offer transferable guidance for future supply-chain information system projects.

4. Problem identification and motivation

The problem identification that follows is based on and structured by the findings resulting from the SLR, fulfilling three interrelated purposes: (1) isolating recurring procurement problems (PPs) observed across industries, (2) examining potential blockchain solution (BS) approaches proposed in the literature, and (3) identifying blockchain-specific implementation problems (BPs). This triadic framing anchors our problem identification firmly in the landscape of existing scholarly solutions.

Procurement Problem Space: The analysis of the procurement problem space was rooted in different sectors, such as the public sector [64–66], construction [51,67,68], oil and gas sector [69], as well as healthcare sector, among others [49,70]. Although the specific challenges differ slightly by sector, the review revealed three overarching categories that can be grouped into: **media discontinuity, process inefficiencies,** and **trust deficit**, as shown in Table 2.

First, **media discontinuity (PP1)** is identified as one of the key procurement problems in 20 of the 34 analyzed papers. The term refers to the lack of seamless data flow between systems. Procurement typically involves multiple stakeholders who rely on diverse tools such as accounting software, business process management systems, and ERP platforms. These systems are often not interoperable across organizational boundaries, leading to fragmented data, disrupted information flows, and manual handovers [65,69–71]. To address this issue, companies commonly implement electronic procurement solutions or establish bilateral agreements to align on technical standards. However, both approaches require standardized data interfaces (i.e., APIs), shared data formats, and semantic interoperability, requirements that are often difficult to fulfill, particularly when large firms impose proprietary standards that smaller vendors must adopt [28,72]. Moreover, traditional ERP systems are frequently costly and inflexible, making it difficult to integrate new partners or adapt existing workflows [72]. These limitations perpetuate organizational silos and slow down procurement processes. In response, several papers highlight blockchain as a promising solution to improve interoperability across organizations. By linking existing databases to a shared, tamper-resistant infrastructure accessible to all authorized participants, blockchain could enhance data integration and coordination [64,65,71].

A second major procurement issue identified in 26 of the reviewed papers is **process inefficiency (PP2)**, which often stems from outdated workflows, rigid organizational routines [17], or opaque payment structures [29]. A primary example is the continued reliance on manual, paper-based procedures, for instance, invoices that still need to be printed, scanned, and processed by hand. Such practices are time-consuming, error-prone, and costly [29,65,73,74]. In this context, blockchain can only deliver added value if organizations first commit to a broader digital transformation. This includes, for example, replacing manual data entry with Internet of Things (IoT) enabled data collection [69]. Once these digital foundations are in place, blockchain can contribute to process improvements by enabling data validation through smart contracts, automating the transmission of information to stakeholders [69], and enhancing auditability through transparent and tamper-proof data sharing [75]. Beyond outdated workflows, inefficient payment systems are frequently cited as an additional barrier to procurement efficiency. Industries structured in cascade-like arrangements with numerous subcontractors and contractual dependencies, exemplarily found in the construction sector, often face delays or errors in payments due to system complexity or opportunistic behavior [29,67,76,77]. Blockchain-based smart contracts offer a potential remedy by enabling funds to be held in escrow and released automatically

Table 1
Overview of the interviewed experts.

Expert	Industry	Role	Experience	Expertise Domain	Duration
E1	Research	Research Assistant	3+ years	Applied Cryptography and Supply-Chain Economics	82 min
E2	Utilities, Research	Founder	5+ years	Energy-Web3 and IoT Innovation	80 min
E3	Supply Chain Finance	Manager	6+ years	DLT Infrastructure and Supply-Chain Solutions	58 min
E4	Supply Chain Finance	Business Dev. Lead	4+ years	Tokenised Deposits and Corporate-Banking Innovation	75 min
E5	Supply Chain Finance	Software Developer	4+ years	DLT Software Development and Supply-Chain Tokenisation	72 min
E6	Research	Research Assistant	5+ years	Blockchain and ZKP Research	99 min
E7	B2B manufacturing	Business process manager	8+ years	Tokenised Business Solutions	89 min
E8	B2B manufacturing	Digital change manager	2+ years	Tokenisation and Supply-Chain Innovation	69 min
E9	Finance	Business Expert	3+ years	Corporate Wholesale-CBDC and Stablecoin	125 min
E10	Research	Research Associate	2+ years	Blockchain and ZKP Engineering	78 min
E11	Research	Professor	7+ years	Blockchain and Cryptography Research	62 min
E12	B2B manufacturing	ERP Change Manager	6+ years	Corporate digital transformation	62 min
E13	Research	Research Associate	2+ years	Information Systems Engineering	99 min
E14	Energy	Procurement manager	5+ years	Large-Scale Energy-Infrastructure Procurement	68 min
E15	Finance	Director	3+ years	Blockchain-based markets analysis	88 min
E16	Consulting, Supply Chain Finance	Project Lead	9+ years	Financial-End-to-End Digital-Asset Value Chain Consulting	72 min

upon fulfillment of predefined conditions. However, to make such mechanisms viable in real-world applications, stable-value tokens are essential, as the volatility of cryptocurrencies makes them unsuitable for long-term contractual payment obligations [28,71,73,76].

Third, **trust deficit (PP3)** is addressed in the vast majority of the analyzed papers (31 of 34), as shown in Table 2. This challenge affects procurement processes across sectors and is rooted in a lack of transparency, accountability, and verifiability. Numerous studies point to fraud risks in various settings, including manipulated tendering procedures in construction and public procurement, favoritism based on personal relationships, and corruption [28,51,66,78]. Blockchain is frequently proposed as a countermeasure to these trust-related issues. By ensuring transparent and tamper-proof transactional records, blockchain can support evidence-based dispute resolution and reduce opportunities for manipulation [17,26,51,73]. The deployment of smart contracts with embedded control mechanisms can further strengthen these benefits.

Blockchain Solution and Problem Space: While the identified procurement challenges and proposed blockchain-based solutions are valid, their effectiveness also depends on how well core blockchain-specific problems, such as the trilemma and privacy, are addressed. To better understand how these aspects were considered, we analyzed which **type of blockchain was used as a solution (BS)** in the reviewed papers and to what extent the scholars engaged with these underlying challenges, as depicted in Table 2. Since not all papers included the development of prototypes to demonstrate their concepts, we further categorized the publications based on the nature of their contributions. Papers that explore how blockchain could address procurement challenges from a conceptual perspective are grouped under theory-oriented contributions, while those involving the development and testing of prototypes are classified as practice-oriented contributions. Observably, researchers tend to prefer Hyperledger Fabric, a private-permissioned blockchain, in situations where privacy or scalability are key requirements. In contrast, Ethereum is more commonly selected when transparency is the primary concern, especially in public sector use cases.

While the blockchain trilemma is rarely named explicitly, many papers briefly touch on individual dimensions of it. Among these is **scalability (BP1)**, which is often treated superficially and was identified as an issue that requires deeper consideration. Some studies implicitly assume that scalability is addressed through the choice of blockchain, such as Hyperledger Fabric, that involves a limited and controlled participant network [70,76]. Others mention scalability as an open issue to be addressed in future research [27,72,74]. Overall, only about one-third of the papers include a more concrete evaluation, such as testing the scalability of their prototypes or proposing technical measures, such as off-chain data storage, to enhance performance [57,80,81].

Similarly, **decentralization (BP2)** is also often addressed only superficially, with many studies treating it as an inherent characteristic of blockchain technology without further discussion [28,65,68,83]. However, relying on a decentralized infrastructure logic does not automatically translate into a truly decentralized implementation in practice. Prototypes are often described as decentralized even when they are built on private-permissioned blockchains that are typically controlled by a single organization or a small group of actors with predefined attributes [27,76]. This significantly limits the actual degree of decentralization achieved [57]. Only a few papers explicitly acknowledge that decentralization depends on the structure of the network and the number of participating nodes, both of which directly influence the system's security and resilience [29,57,71,82].

Table 2
Literature overview.

Literature		Procurement Problem Space			Blockchain Solution and Problem Space				
Type	Item	Media discontinuity (PP1)	Process inefficiency (PP2)	Trust deficit (PP3)	Blockchain type (BS)	Scalability (BP1)	Decentralization (PB2)	Security (BP3)	Privacy (BP4)
Practice Oriented Contribution	[70]	x	x	x	ETH	(x)	(x)	(x)	–
	[65]	x	x	x	HLF	x	(x)	(x)	x
	[79]	x	x	x	ETH	–	–	–	x
	[66]	–	–	x	ETH	–	–	–	–
	[51]	–	–	x	ETH	–	(x)	(x)	–
	[78]	–	x	x	ETH	–	(x)	(x)	x
	[67]	–	x	x	ETH	–	(x)	(x)	(x)
	[68]	x	x	x	HLF	–	(x)	(x)	x
	[57]	–	x	x	HLF	x	x	x	x
	[77]	x	x	x	HLF	x	–	(x)	x
	[80]	–	–	x	BT	x	(x)	(x)	x
	[81]	–	x	x	ETH	x	(x)	x	x
	[82]	x	x	x	pETH	x	x	(x)	x
	[74]	x	x	–	HLF	(x)	(x)	(x)	x
	[27]	x	x	x	HLF	(x)	x	(x)	x
	[43]	–	x	x	Layers	x	(x)	(x)	x
	[83]	–	x	x	ETH	x	(x)	(x)	x
	[76]	–	x	x	HLF	(x)	(x)	x	x
	[29]	x	x	x	ETH & HLF	(x)	x	x	x
Theory Oriented Contribution	[52]	–	–	x	PND	–	–	(x)	(x)
	[84]	x	x	x	PND	(x)	–	–	–
	[71]	x	x	x	various	(x)	(x)	(x)	(x)
	[85]	x	x	x	–	–	(x)	(x)	–
	[75]	x	x	x	–	–	(x)	(x)	–
	[17]	x	x	x	PND	(x)	(x)	(x)	x
	[49]	–	x	x	PND	(x)	(x)	x	–
	[73]	x	x	x	–	x	(x)	(x)	(x)
	[86]	x	–	x	–	(x)	(x)	(x)	–
	[26]	–	x	x	–	(x)	(x)	(x)	(x)
	[87]	–	–	x	–	x	(x)	x	(x)
	[72]	x	–	–	–	(x)	(x)	(x)	(x)
	[28]	x	x	x	HLF	(x)	(x)	(x)	(x)
	[64]	x	x	x	PND	(x)	(x)	(x)	(x)
	[69]	x	x	–	PND	(x)	(x)	(x)	(x)

x : addressed; **(x)** : partially addressed; – : not addressed.
ETH : Ethereum; **HLF** : Hyperledger Fabric; **BT**: Bitcoin **PND** : Permissioned.

A similar observation applies to **security (BP3)**, which is often assumed to be an inherent feature of blockchain technology. However, some papers do address potential vulnerabilities, such as single points of failure in private-permissioned blockchains and various attack vectors that can compromise blockchain systems or smart contracts [28,67,81]. These include risks such as Denial-of-Service (DoS) attacks, which can disrupt operations by overwhelming the network, and other targeted exploits that must be considered when designing blockchain-based solutions.

In contrast to the blockchain trilemma, **privacy (BP4)** is discussed more thoroughly, especially in the papers with practical contributions. While blockchain's transparency is often seen as beneficial, many scholars acknowledge the risk of exposing sensitive data and therefore propose solutions such as data encryption, off-chain storage, or the use of private-permissioned blockchains with private channels to restrict data visibility to selected participants [28,64,65,78,80]. Accordingly, blockchain implementations in procurement should recognize the danger of data exposure as a fourth, overarching challenge.

Based on the findings of the literature review, we distilled a concise problem statement to steer our subsequent design work: While blockchain is widely recognized as a promising remedy for procurement inefficiencies, the solutions documented to date overwhelmingly favor private or consortium architectures that achieve scalability and confidentiality by sacrificing decentralization, which is often coupled with security. Hence, broad real-world adoption still depends on tackling foundational design tensions, most notably the blockchain trilemma and stringent privacy requirements, which current approaches only partly resolve. In particular, the literature reveals a gap regarding the joint recognition of the trilemma and privacy, and the possibility and benefits of addressing it by combining public blockchain infrastructures with emerging tools such as ZKPs and rollups. The next section, thus, shows how the evaluated DOs were derived, which respond both to the procurement challenges identified earlier and to the architectural prerequisites needed to unlock blockchain's full potential.

5. Evaluated design objectives

Following the approach proposed by Peffers et al. (2007) we translated the previously identified procurement challenges and blockchain-related design tensions into DOs for our artifact, which were then iteratively adapted and evaluated in the course of our qualitative interview cycle [36]. In the following, we present the final set of five evaluated DOs that guide our design process, explain their rationale identified in the literature, and incorporate the practical insights gained from the expert interviews. Table 3 offers a concise overview of each DO, its description, and its empirical grounding in both the literature and the evidence from interview data.

As outlined in the previous section, media discontinuity remains a common challenge across industries [28,65,69–72]. Interviewees confirmed this with concrete examples, such as disruptions caused by partner system changes (E14) or the lack of synchronized data across multiple ERP systems within a single company (E4). While all interviewed experts agreed that standardization is key to addressing this issue, they also emphasized the difficulty of reaching consensus on a shared standard. Private blockchains can enhance interoperability, but they typically require participants to join a consortium and adopt a

Table 3
Evaluated Design Objectives.

DO	Evidence from Literature	Description	Evidence from Expert Interviews
DO1	PP1, BP2	The artifact should leverage standardization efforts to enable enterprises to seamlessly integrate their ERP systems with publicly accessible, interoperable blockchain infrastructure.	E4, E14, E15, E16
DO2	PP2	The artifact should embed stable and regulation-compliant payment processing into blockchain-based procurement workflows, enabling efficient and seamless value exchange among participants in both national and international contexts.	E4, E6, E7, E9, E15
DO3	PP3, BP2, BP3	The artifact should strengthen trust in procurement's dispute resolution processes without undermining trust in the underlying infrastructure by leveraging permissionless blockchain.	E3, E4, E5, E6, E7, E8, E11, E12
DO4	BP1	The artifact should be built on an infrastructure capable of scaling with growing transaction volumes and users.	E1, E2, E5, E6, E10
DO5	BP4	The artifact should be privacy-enabled while maintaining the inherent auditability of blockchain records.	E2, E4, E7, E8, E16

predefined standard. As E15 noted, *"Creating a consortium-based standard might become problematic when trying to integrate more customers"*. E16 added that although private blockchains may offer short-term benefits in establishing common standards, it believes that only public solutions are likely to be sustainable in the long run, having observed multiple consortia, even those with few participants, fail due to internal disagreements on standards. These insights point to the need for a more openly accessible approach. Consequently, we define **DO1** as following: **The artifact should leverage standardization efforts to enable enterprises to seamlessly integrate their ERP systems with publicly accessible, interoperable blockchain infrastructure.** This approach reduces the need to join consortia or engage in complex negotiations. Thus, a design based on public blockchain infrastructure creates a neutral and open environment, lowering adoption barriers and fostering a competitive landscape in which effective standards can emerge through practical implementation and broad acceptance.

Furthermore, as outlined in the previous section, procurement processes are affected by various inefficiencies, including manual interventions [29,65,74], as well as payment-related inefficiencies caused by complex cascading structures and opportunistic behavior [67,76,77]. As noted by several interviewees (E4, E6, E7), payment-related challenges vary significantly across industries and regions. For instance, issues that were previously widespread in the EU, as described by E15, have largely been mitigated through the introduction of instant payment regulations. However, inefficiencies persist in sectors such as construction [74], and in cross-border transactions, as highlighted by E9, where settlement processes involve multiple parties and intermediary banks, leading to delays. Consequently, while the literature [69] and expert E7 emphasize that many process-related inefficiencies should be addressed through conventional digitization before introducing blockchain, payment inefficiencies are of a different nature. These challenges often do not stem from a lack of digitization, but rather from the involvement of multiple stakeholders and the resulting coordination complexity. Blockchain technology can help mitigate these issues by providing a shared, transparent infrastructure that enables automated, traceable, and rule-based value transfers. In multi-party settings, it can significantly reduce delays, improve accountability, and enhance overall process efficiency. For this reason, we chose to focus primarily on payment-related inefficiencies rather than general process inefficiencies that require digitization first. Furthermore, as emphasized by E4, *"separating payment execution from blockchain-based procurement systems creates isolated data environments that hinder process integration"*. To be viable in practice, embedded payment functions must also comply with regulatory requirements to ensure legal compatibility. Hence, to address these challenges, **DO2** requires that **the artifact should embed stable and regulation-compliant payment processing into blockchain-based procurement workflows, enabling efficient and seamless value exchange among participants in both national and international contexts.**

While the trust deficit highlighted in the literature [28,51,66] was confirmed by the interviews, concerns specifically related to fraud and corruption were perceived as less severe within the industrial context and are more applicable to the public sector. Several industry experts noted that they are less worried about being defrauded by trading partners and more concerned about the integrity and reliability of their partners' internal processes (E3, E4, E7, E8). Manual steps that still dominate procurement processes frequently lead to errors, which are time-consuming and costly to resolve (E4, E5). Researchers supported this view by pointing out that the oracle problem persists across all infrastructures, centralized and decentralized alike (E6, E7, E11). As a result, trust can be meaningfully enhanced only by strengthening the mechanisms for dispute resolution, making them faster, more traceable, and ensuring that parties are held accountable for the data they provide (E4). Blockchain infrastructure, therefore, appears particularly suitable for increasing procedural trust, rather than trust in the parties themselves. Conventionally, companies rely on centralized solutions such as electronic procurement platforms, which act as third parties in dispute resolution. However, as noted by E5 and E12, centralized systems, including private blockchains, introduce single points of failure and require thorough security assessments, as they can undermine trust in the resolution process. In the blockchain domain, decentralized infrastructure, particularly permissionless systems that allow open participation in consensus, is widely assumed to offer the highest level of security, as emphasized by E6, E8, and E12. Consequently, **DO3** requires that **the artifact should strengthen trust in procurement's dispute resolution processes without undermining trust in the underlying infrastructure by leveraging permissionless blockchain.**

Since the first and third DO jointly address procurement challenges as well as the decentralization and security dimensions of the blockchain trilemma, by requiring the use of a public-permissionless infrastructure that benefits procurement solutions, we do not explicitly cover these dimensions again in separate DOs. Nevertheless, the security of the smart contract logic must still be assessed independently in each individual case [57]. While decentralization and security enhance trust and resilience, they are often linked to limited scalability, a factor

that is central to system functionality. This trade-off, as described in the blockchain trilemma, is rarely sufficiently addressed by procurement artifacts built on blockchains [51,67,78]. Meanwhile, several interviewed experts (E1, E2, E10) noted that scalability remains a concern not only for public-permissionless blockchains but also for private-permissioned infrastructures, particularly as the number of users grows. In such cases, the system must manage increased computational redundancy. Moreover, some clients sometimes expect the provider of the blockchain infrastructure to operate their nodes as a service (E5, E6), further increasing complexity and scalability demands. As a result, and in line with all interviewed experts, we defined in **DO4** that regardless of the blockchain type, **the artifact should be built on an infrastructure capable of scaling with growing transaction volumes and users**. This includes maintaining low latency, high processing capacity, and stable transaction costs, even under increased system load.

Moreover, while blockchains' inherent transparency can help to build trust, it also poses a threat to sensitive data, as highlighted in our analyzed literature [28,65,78,80] and confirmed during the interview cycle. Several experts (E2, E4, E7, E8, E16) emphasized that data should be accessible on a need-to-know basis, either by protecting sensitive information through selective disclosure or by storing it entirely off-chain, with relevant data made viewable only to authorized parties for traceability purposes. This ensures that the benefits of blockchain are preserved. Consequently, in **DO5**, we define that **the artifact should be privacy-enabled while maintaining the inherent auditability of blockchain records**. It should offer users the flexibility to protect business-critical and sensitive information from unauthorized access, while still ensuring traceability for relevant procurement stakeholders and regulatory authorities as demanded by the second and third DO.

6. Design and development of the artifact

As outlined in Section 3, after analyzing various existing technologies, we decided to apply the Aztec privacy-oriented L2 protocol in order to fulfill predefined DOs. The following subsection summarizes Aztec's architecture, followed by a detailed discussion of the operational logic of the resulting smart-contract ecosystem. We focus the following description on the architectural role and organizational implications of the technical components, rather than on cryptographic implementation details. In doing so, we incorporate a physical view in Fig. 4 as well as a process view of our architecture in Figs. 5 and 7, following the guidelines of Kruchten (2002) [37].

6.1. Infrastructure and components of the artifact

To understand how the procurement smart contract ecosystem operates within Aztec, it is first necessary to grasp the fundamentals of its L2 architecture. Built on top of Ethereum L1, Aztec inherits Ethereum's robust security and data availability guarantees, thereby directly supporting DO1 and DO3 by ensuring public accessibility and leveraging a permissionless base to foster trust in the infrastructure [35]. The connection between the layers is facilitated through a decentralized communication network and Aztec's APIs and libraries, which help establish a publicly accessible integration standard, further reinforcing DO1. However, Ethereum's base layer makes all transaction data publicly visible, even if addresses remain pseudonymous. To mitigate this exposure, Aztec applies ZKPs coupled with additional components explained further below, thus delivering confidential, high-throughput transactions while retaining decentralization. The protocol's core components directly advance DO4 (scalability) and DO5 (privacy) for multi-party procurement scenarios, which will be described in the following sections.

First, to guarantee privacy of input data (DO5), the execution of private functions is moved off-chain and is done locally by users within a Private Execution Environment (PXE). In this environment, users' devices generate ZKPs that attest to the correctness of their actions, acting as the prover that convinces another party (the verifier) that a statement is true without disclosing the underlying data [31]. Aztec specifically utilizes Zero-Knowledge Succinct Non-Interactive Arguments of Knowledge (zk-SNARKs) and implements the PLONK protocol, a universal and efficient zk-SNARKs system [88]. PLONK and its different extensions, such as Turbo PLONK and UltraHonk, enable fast proof generation, efficient verification, and reuse of a single trusted setup across different circuits in the Aztec network, making it especially well-suited for general-purpose, private smart contracts on public blockchains (DO4, DO5). While the underlying cryptographic protocols are technically complex, their functional role in the artifact can be summarized as follows: they enable participants to prove that procurement transactions comply with agreed rules (e.g., authorization, balance sufficiency, and contractual conditions) without revealing prices, quantities, or counterparty relationships to other network participants.

Second, these ZKPs are verified and finalized on Ethereum L1 through a decentralized communication network among layers, such as the sequencer network (i.e., a set of nodes responsible for ordering and batching L2 transactions), which are free to participate by staking on Ethereum, thus supporting privacy without compromising DO3's demand for permissionlessness. After verifying each proof from PXEs, the sequencer bundles private and public transactions into a single proof and submits it to the Ethereum network for verification and finality. This architecture enables scalable data handling without compromising Ethereum's base-layer decentralized security (DO3) and constructed privacy (DO5). Final settlement occurs through a rollup smart contract on Ethereum, which stores Merkle roots of state trees to ensure the integrity of contract logic and private/public data states [89].

Third, the communication between L1 and L2 on Aztec is facilitated through secure message bridges, known as portals [90]. Portals are smart contracts that enable the exchange of encrypted messages between layers while preserving confidentiality. One portal always resides on L1, while its counterpart operates within the L2 system. Aztec has also deployed Inbox and Outbox smart contracts on L1 to support efficient data transmission. These contracts collect incoming and outgoing messages, allowing sequencers to retrieve and process them. Messages are thereby transmitted as cryptographic hashes, and a nullifier scheme to ensure that only the intended recipient, who possesses the corresponding secret, can consume the message on L2. This architecture allows Aztec to offer strong privacy guarantees without compromising Ethereum's trust assumptions.

Although the Aztec protocol encompasses complexities beyond the scope of this brief description, the overview presented here establishes the architectural foundation for our prototype, whose principal components are also depicted in Fig. 4. The left side of the figure represents smart contracts written in Solidity and deployed on the public layer. At the core of the accounts is the concept of Account Abstraction, where users interact through smart contract accounts rather than traditional Externally Owned Accounts (EOAs) [91]. This enables businesses to implement custom access controls, sponsor transaction fees, and improve security with features like multi-signature approvals and key recovery. In the context of procurement, the two accounts depicted in Fig. 4 represent two companies wishing to engage in procurement. The Token smart contract, deployed by a Financial Institution (FI), mints and distributes ERC-20 stablecoins to customers. By supplying participants with a value-stable and regulatory-compliant medium of exchange, the contract directly supports DO2. The Token Portal smart contract is responsible for creating encrypted messages transmitted between layers. Furthermore, the green-colored Outbox and Inbox smart contracts in Fig. 4 facilitate the communication between L1 and L2, while the rollup is mainly used to keep track of state roots, perform state transitions, and ensure data availability.

The right side of Fig. 4 provides a simplified view of the key L2 components developed for this system. All smart contracts on this layer are written in Noir, a domain-specific language tailored for zero-knowledge applications on the Aztec network [92]. Noir allows for automatic ZKPs

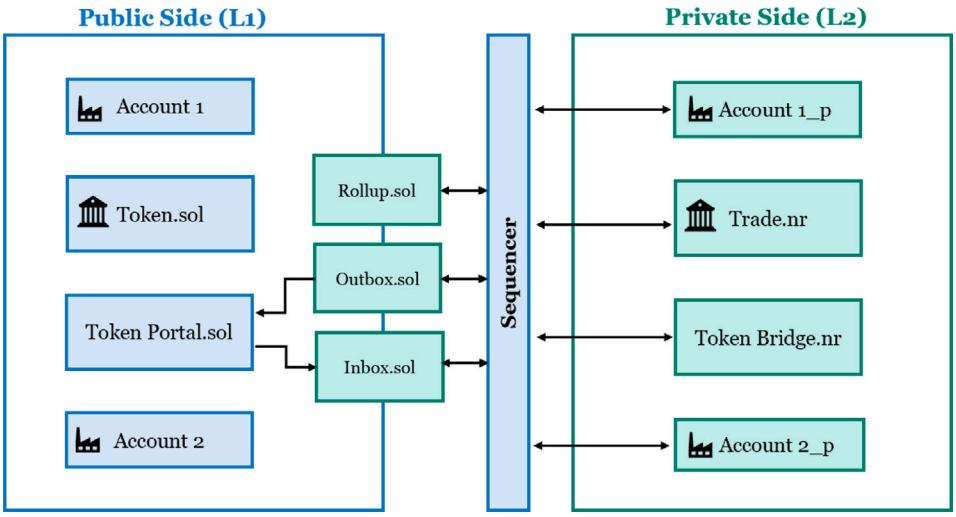

Fig. 4. Overview of components and communication flow between layers.

generation and leverages Aztec's modular architecture, thus lowering integration barriers and accelerating adoption through reusable, well-documented components and simplified ZKPs generation. The accounts on L2 replicate the functionality of publicly visible L1 accounts while ensuring complete confidentiality. The Token Bridge smart contract, which corresponds to the Token Portal on L1, is responsible for receiving messages and making them accessible to participants on L2. Finally, the Trade Contract, supposedly deployed by a FI to facilitate the required regulatory checks (DO2), orchestrates the entire procurement workflow on the L2 network, from supplier selection to final payment, thereby delivering a secure and privacy-preserving end-to-end process.

6.2. Exemplary minting and transfer of tokens from L1 to L2

In the following, we will illustrate the artifacts' operation through representative examples. To test and demonstrate the functionality of the conceptual artifact and the corresponding examples, we used the Aztec Sandbox (v.24), as no public testnet was available at the time of the development. The sandbox was deployed in a Docker-based environment that included Anvil for simulating a local Ethereum instance and a local Aztec rollup to emulate L2 operations. Smart contracts were developed in Solidity for L1 and Noir for L2. Testing was conducted using TypeScript, allowing programmatic interaction with the deployed contracts, testing different phases, and identifying possible functional errors.

As the background section outlines, procure-to-pay activities comprise five phases: need identification, supplier selection, ordering, fulfillment, and settlement. Regulatory requirements are present across all phases, yet they become most acute during settlement, where obligations related to payment, custody, and statutory reporting converge and must be addressed comprehensively. Consequently, we first focus on how procurement parties get access to regulatory compliant tokes for later settlement. In the simplified scenario used for our prototype demonstration and expert interviews, these compliance tasks are exemplarily handled by a FI that is licensed under the Markets in Crypto-Assets Regulation (MiCAR) to issue and distribute stablecoins [93]. While concentrating the role of the FI in a single institution may appear to create a tension between public accessibility (DO1) and regulated processes (DO2), this is primarily a consequence of the deliberately minimal prototype setting. From an economic perspective, an FI has strong incentives to maximize participation, since the value of the system increases with adoption due to network effects, rather than to restrict access. Moreover, participants that are admitted by an FI do not need to operate their own nodes or set up dedicated infrastructure, which keeps the barrier to entry low. The purpose of this design choice is to validate regulatory compliance in a controlled setting, and the architecture can be straightforwardly extended to multi-institution scenarios, which would further mitigate this tension.

Furthermore, although the prototype illustrates regulatory compliance using MiCAR as a concrete reference framework, the underlying architectural design is not limited to a single jurisdiction. In particular, data-protection regimes such as the General Data Protection Regulation (GDPR) impose additional constraints on the handling of personal and commercially sensitive information in inter-organizational and cross-border procurement processes. The proposed layered architecture supports such requirements through data minimization and selective disclosure principles. Sensitive business and transactional data are processed off-chain in encrypted form, while the public base layer stores only cryptographic commitments, thereby avoiding the persistence of personal or confidential information on immutable ledgers. This separation aligns with core GDPR principles, such as purpose limitation and data minimization, and facilitates compliance across jurisdictions without compromising the system's trust and auditability guarantees.

More broadly, cross-border procurement regulations differ not only in reporting obligations and regulatory access rights, but also in data localization requirements and the degree of supervisory intervention. The design accommodates such heterogeneity by modularizing compliance-related functionality, such as whitelisting logic, regulatory visibility mechanisms, data retention policies, and access control rules, at the application and L2 levels. This modular separation enables jurisdiction-specific adaptations, ranging from stricter data sovereignty regimes to more centralized oversight models, while preserving the permissionless trust anchor of the base layer.

The process starts when the company completes Know-Your-Customer (KYC) and other due diligence checks with the FI, as outlined in Fig. 5. After approval, the company deposits fiat currency at the account of FI, which triggers the issuance of an equivalent amount of stablecoins. Following the off-chain fiat deposit, the FI may transfer the stablecoins to the company's existing L1 blockchain address or create a new account on its behalf. The FI whitelists the firm's public L1 address (A1) and its corresponding private L2 address (A1_p). The Token contract then mints the stablecoins to A1, where they can circulate on L1 before being moved to A1_p for confidential procure-to-pay transactions.

If a company prefers complete privacy, it may request the direct issuance of stablecoins into a private L2 account. However, the presented example assumes that account A1 wishes to hold one part of its tokens publicly for alternative exchanges and conceal another part by transferring them into a private environment. This might be particularly relevant in procurement contexts, where, e.g., traded prices

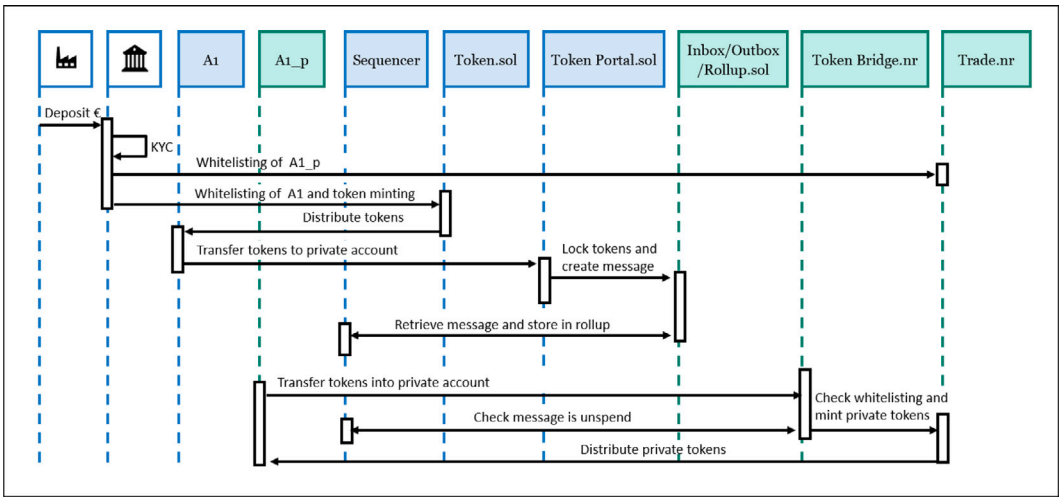

Fig. 5. Processual view on token transfer from L1 to L2.

should remain confidential and inaccessible to competitors. To initiate the transfer to L2, A1 interacts with the Token Portal smart contract, which locks the specified token amount and generates a corresponding transfer message. Given the public nature of L1, no private address must be disclosed during this transfer. To preserve privacy, A1 submits the transfer amount and a hashed secret as input parameters. Hashing is a cryptographic technique that transforms input data (in this case, a secret known only to A1) into a fixed-size string of characters [44]. This string, or hash, acts as a unique fingerprint of the original input but does not reveal it. In this context, the hashed secret serves as a commitment, which can later be shown and verified on L2 to prove ownership of the transferred tokens without exposing sensitive information. The Token Portal then forwards the transfer message to the Inbox smart contract, where a sequencer within the decentralized network picks it up. The sequencer processes the message and incorporates it into a message Merkle tree, the root of which is stored in the Rollup smart contract on L1 [89].

Subsequently, the company uses its private L2 address (A1_p) to request the token distribution via the Token Bridge smart contract on L2. To do so, the user submits the transfer amount and the unhashed secret, enabling the Token Bridge to compute the hash and verify the message's existence and unused status via the sequencer. Upon successful verification, the Token Bridge invokes the Trade smart contract, which mints and distributes the private tokens to the user's L2 address. A final whitelisting check is performed during this step to ensure that the FI retains regulatory oversight, even in the privacy-preserving environment of L2. In this case, both the identity of the account holder and the transferred amount are concealed.

In either scenario, account addresses must be whitelisted by the FI to enable ongoing monitoring and regulatory oversight. This functionality is enforced through smart contracts: the Token Contract on L1 and the Trade Contract on L2, both of which are deployed by the FI to ensure regulatory compliance across layers. In our simplified setup involving a single FI, the Trade Contract governs private token issuance and transfers on L2. However, introducing multiple financial institutions and various token issuers would necessitate re-evaluating the whitelisting and compliance mechanisms to maintain regulatory consistency across a more complex ecosystem. To prove the functionality the test results in Fig. 6 under Test B illustrate the transfer of 10,000 tokens by A1 to its private address (A1_p). While viewable for test purposes, the balance information would remain confidential in production environment and can only be retrieved by the respective token holders, thereby preserving the system's privacy guarantees.

6.3. Exemplary privacy preserving procure-to-pay process on L2

After detailing the mechanisms that enable token transfers between L1 and L2, the following section demonstrates how the artefact supports a compliant, privacy-preserving procure-to-pay process, using a simplified representation of the process's typical steps. To illustrate the core functionality of the developed artifact in a clear, focused manner suitable for the expert interviews, we model a minimal scenario that involves only two companies: a supplier (A1_p) and a buyer (A2_p). Both firms have deposited Euros into an FI account, received private tokens for trading, and have been identified and whitelisted by the FI for regulatory compliance. Fig. 7 depicts the resulting exemplar procure-to-pay process. Although real-world procure-to-pay workflows are far more complex and engage multiple stakeholders, systems, and interdependent sub-processes, this streamlined example suffices to demonstrate the artifact's key capabilities.

The initial step of a procure-to-pay process is need identification, which typically occurs within a company and may be triggered by various factors, such as the need to replenish inventory, procure new production inputs, or address unexpected operational requirements. This process happens off-chain and is beyond the present scope. Nevertheless, with further digitization of internal procedures, this step could be integrated into a broader end-to-end automation initiative. Once a requirement has been identified, the company proceeds to supplier selection, based on predefined criteria such as price, quality, lead time, regulatory compliance, and reliability [1]. To enable this interaction in the developed artifact, the first step is to make the available products visible to potential counterparties. As shown in Fig. 7, A1_p must first register a product by providing its identifier, optional descriptive attributes, and public key, allowing potential trading partners to identify and contact the supplier. This information is stored in a publicly accessible list. Subsequently, A2_p, upon identifying its internal need, can consult this list to find a matching product and retrieve the supplier's public key.

A2_p subsequently advances to the ordering phase by invoking the Demand function, submitting a request that specifies quantity, price, delivery date, and, where relevant, additional quality requirements. This part of the procure-to-pay process must be executed privately. Consequently, private functions facilitate localized execution where raw inputs remain on the user's device, while the resulting output is transformed into an encrypted private log and a Zero-Knowledge Proof (ZKP) is generated to verify the transaction's validity. The sequencer then collects, verifies, and aggregates these proofs into a rollup batch. For permanent on chain integrity, a cryptographic commitment in the form of a hash is posted to L1, whereas the encrypted log is stored

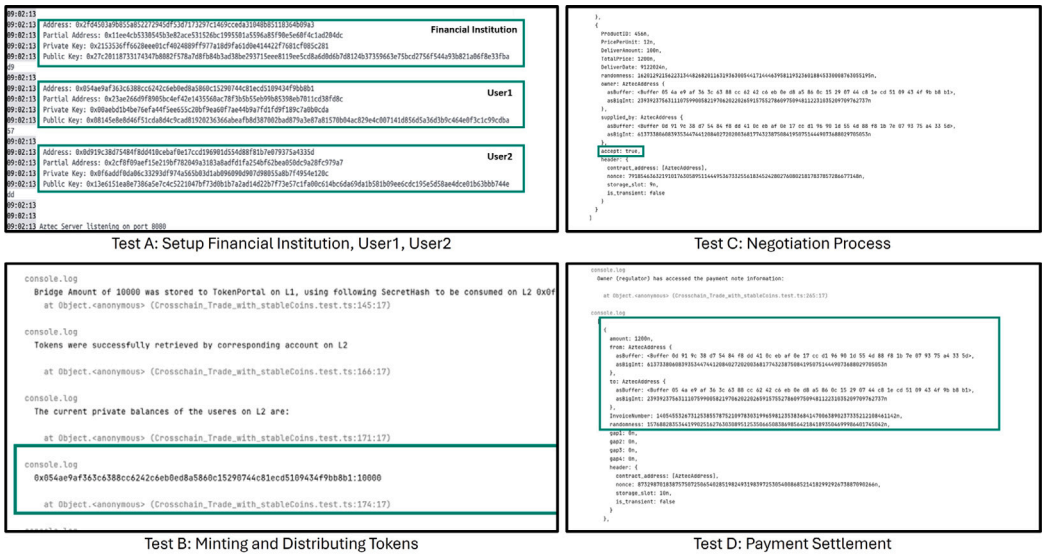

Fig. 6. Console view on the different test scenarios.

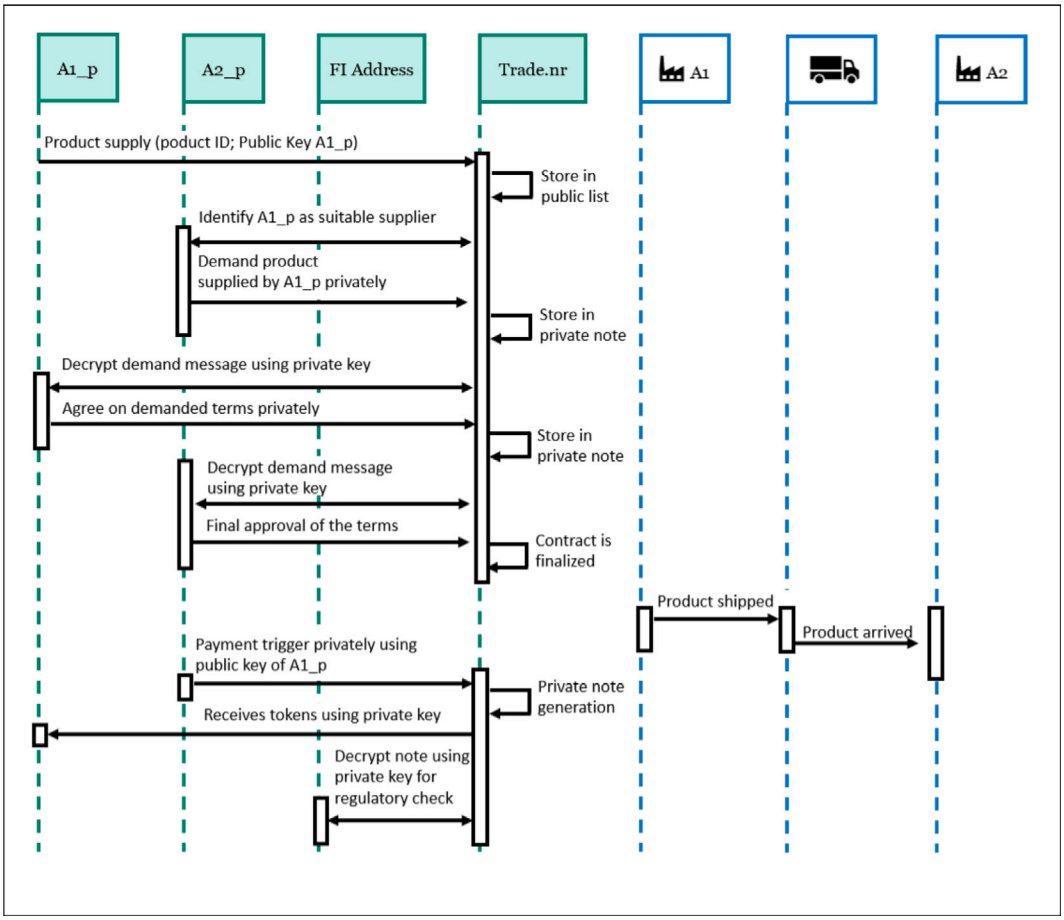

Fig. 7. Procure-to-pay process.

within a blob, a cost efficient, temporary data container, for a limited period. This window allows the counterparty's (A1_p) PXE to continuously scan and decrypt the information before the blob is pruned from the network nodes. This architecture ensures that sensitive trade data remains confidential and storage overhead is minimized, while the persistent hashing on the mainnet guarantees immutable traceability.

A1_p can then respond in the same encrypted manner, using the public key of A2_p to send an offer that includes price, quantity, and other relevant terms. This private exchange can continue iteratively until both parties reach an agreement. The contract is considered finalized once the demand and offer notes match exactly. At this point, an order and invoice are automatically generated and made accessible only to

the relevant parties. Depending on sector-specific needs, the involved companies may also confidentially exchange further information such as company names, dispute resolution rules, or other contractual details, ensuring these remain hidden from competitors and unauthorized third parties.

To showcase the functionality, Fig. 6 under Test C displays the decrypted notes of A2_p after negotiating the agreement terms for Product456 with A1_p. The decrypted note contains information accessible only to A2_p, including terms such as price and delivery date, the involved public keys, and additional randomness for enhanced security. The accept: true flag indicates the agreement status, confirming that both parties have entered into a contract, with the outcome documented in an immutable manner.

The next phase, fulfillment, occurs off-chain. In this step, A1 arranges a logistics provider to deliver the ordered goods to A2. All logistics-related operations are assumed to take place outside the blockchain environment. However, since the order and invoice information can be recorded on-chain, dispute resolution is significantly simplified and accelerated, serving the DO3. After delivery, A2_p manually initiates payment, entering the settlement phase. Because the scenario lacks IoT-based proof-of-delivery, payment relies on buyer confirmation, mirroring industry practices where terms often extend up to 60 days. A2_p initiates the transaction using a designated function within the Trade smart contract. Nevertheless, the programmable nature of the artifact allows for future automation, linking payment execution to predefined conditions such as delivery confirmations or sensor data.

Finally, as outlined at the beginning of the previous section, settlement falls under regulated activities. Consequently, each payment must undergo verification by the corresponding FI to ensure compliance with Anti-Money Laundering (AML) and Countering the Financing of Terrorism (CFT) regulations (DO2). For this purpose, the transfer function within the trade smart contract includes an embedded feature that stores the public key of the FI's private account. This enables the FI to decrypt transaction details using its private key while keeping the data confidential for all other parties, thus ensuring compliance with AML obligations, which require the FI to retain such records. Additionally, the built-in whitelisting mechanism verifies that the sending and receiving parties are not subject to sanctions, ensuring compliance with CFT requirements.

Fig. 6 under Test D exemplarily demonstrates the note accessible to the FI after the transaction in the test environment. In the constructed example, A2_p transferred 1200 tokens to A1_p. The note contains critical information necessary for regulatory oversight, including the transferred amount, the sender and receiver addresses, and a corresponding invoice number that serves as a reference for identifying the purpose of the transaction. This design ensures that regulatory bodies can perform compliance checks without compromising the privacy of the involved parties vis-à-vis other network participants.

The artifact delivers a privacy-preserving, regulatory-compliant infrastructure that successfully supports a simplified yet fully functional procure-to-pay workflow. Covering the core process steps establishes a robust foundation for subsequent automation and integration efforts. As additional intra- and inter-organizational processes become digitized, capabilities such as automated need detection, dynamic supplier matching, and condition-based settlement can be added seamlessly, enhancing efficiency, transparency, and trust across digital procurement ecosystems.

7. Evaluation

This section reports the evaluation of the developed artifact in accordance with the strategy outlined in the Method section. Following DSR principles, the evaluation combines a quantitative viability assessment with a qualitative expert-based evaluation to address complementary dimensions of the research question. While the quantitative evaluation assesses whether the artifact can plausibly meet operational requirements under realistic procurement workloads, the qualitative evaluation evaluates whether the artifact conceptually and functionally meets the predefined design objectives. The quantitative viability assessment is presented first and focuses on throughput, latency, and economic feasibility. It is explicitly framed as an assessment of operational plausibility rather than definitive performance benchmarking. Given the prototype status of the artifact and the evolving maturity of the underlying L2 infrastructure, this assessment combines controlled sandbox experiments, workload projections, and protocol-level characteristics to assess whether the proposed architecture exhibits feasible scaling behavior for inter-organizational procurement processes. The results of the quantitative viability assessment are summarized in Table 4. The qualitative expert evaluation follows and provides a structured, summative evaluation of the artifact's alignment with the DOs. The implementation code, test results, and conceptual views of the final prototype were used to demonstrate the artifact's functionality to domain experts during qualitative interviews.[1] During these guided walkthroughs, the architecture and workflows described in Section 6 were presented to the experts in the context of concrete procurement scenarios. Verbatim expert feedback was systematically coded and aggregated according to explicit decision rules, distinguishing between full and qualified fulfillment and non-fulfillment of DOs based on recurring assessment patterns. The results of the qualitative evaluation are summarized in Table 5. Filled circles indicate consistent expert agreement that a design objective is sufficiently fulfilled, while half-filled circles indicate qualified fulfillment accompanied by bounded concerns or contextual limitations. Empty circles would indicate non-fulfillment. However, such cases do not appear in the table because indications of non-fulfillment triggered a return to the design and development stage during earlier build–evaluate cycles. The following subsections present the quantitative and qualitative evaluations in detail.

7.1. Quantitative viability assessment

Although our prototype operates within a localized Docker-based sandbox and the Aztec network currently offers limited public data, this evaluation establishes a quantitative framework to assess the system's viability. We analyze three core dimensions necessary to meet the predefined DOs: **Scalability** (throughput), **latency** (time-to-finality), and **economic sustainability** (cost efficiency). By combining experimental results with established estimates, we assess the system's applicability to real-world procurement cases. The results are summarized in Table 4.

First, **scalability** must be evaluated to ensure the system can handle real-world interorganizational procurement volumes, as required by DO4. A reliable proxy for this Business-to-Business (B2B) interaction intensity is invoice volume, as invoices document the actual interorganizational exchange of information about goods and services within procurement processes. According to the European Commission, approximately 18 billion invoices are exchanged annually in the European Union [94]. At the same time, Eurostat reports that there are about 33.1 million active enterprises. Normalizing these figures implies that an average enterprise processes roughly 540 invoices per year, corresponding to about 2.2 invoices per business day, assuming 250 working days [95]. To translate invoice volume into system load, we estimate the number of verifiable events per invoice, as each invoice results from a multi-stage process comprising negotiation, fulfillment, invoicing, and settlement. In our prototype, contract negotiation is also executed on-chain and precedes these execution stages. Assuming a minimal negotiation trace (offer, counter-offer, final offer, acceptance), four negotiation interactions, plus invoicing, fulfillment, and

[1] Code and test results available at: https://doi.org/10.5281/zenodo.17218727.

settlement, yield a conservative estimate of about 7 verifiable events per invoice. Thus, for a mid-sized supply chain network with 1000 participating firms, the normalized activity level corresponds to about 2200 invoices per business day. Treating each invoice as the result of approximately 7 verifiable events implies a daily workload of roughly 15,400 transaction events.

While this number corresponds to an average load far below one Transactions per Second (TPS) when uniformly spread over 24 h, real-world enterprise workloads are concentrated during business hours and at end-of-day processing. Assuming that 80% of daily activity occurs within an 8-hour window, the average load during active hours is about 0.4–0.5 TPS ($(0.8 \cdot 14,500)/(8\text{h} \cdot 3600\ \text{s})$). However, in practice, systems are dimensioned for peak rather than average load [96]. Consequently, we conservatively assume a one- to two-order-of-magnitude headroom. Using the upper bound of the average load of 0.5 TPS as a baseline, a 10× factor yields $0.5 \cdot 10 = 5$ TPS as a sustainable throughput level, while a 100× factor yields $0.5 \cdot 100 = 50$ TPS as an extreme burst scenario. Using the lower bound of the average load of 0.4 TPS together with an intermediate peak factor of 50× yields $0.4 \cdot 50 = 20$ TPS, which we take as a conservative lower bound for peak capacity. Accordingly, we target a system capacity in the range of 20–50 TPS. In the privacy-preserving architecture considered here, the targeted throughput range also reflects the computational overhead associated with ZKPs generation. Because each transaction batch requires the construction of validity proofs, proving time directly affects how many batches can be sequenced, included, and made visible at the L2 level per unit of time, and thus the achievable TPS. Sustained high transaction rates also increase the proofing workload per epoch, which may indirectly affect how quickly aggregated proofs can be submitted and verified on L1. As circuit complexity and the number of private state transitions increase (DO5), longer proving times may reduce effective throughput, particularly during peak activity periods, and should therefore be carefully considered in system design.

To obtain an initial empirical reference point for the scalability of our design, we conducted controlled experiments in the local Docker-based prototype using a single sequencer and executed batches of 10, 20, and 40 transactions, each repeated 3 times. For 10 transactions, execution times ranged from 201 to 220 s (mean $\approx$ 210 s), for 20 transactions from 398 to 415 s (mean $\approx$ 406 s), and for 40 transactions from 580 to 611 s (mean $\approx$ 596 s), corresponding to an effective throughput of approximately 0.05–0.07 TPS in the current setup. Larger batches could not be processed reliably in this environment, as the TypeScript-based orchestration required artificial delays to prevent execution failures, indicating that the observed limitations arise from the single-node Docker-based prototype setup. Likewise, the current public test network Ignition is primarily intended for testing protocol and governance mechanisms, generating proofs for empty blocks rather than for realistic performance benchmarking [97]. However, with over 2000 participating sequencers, the network already provides a strong indication of the intended decentralization (DO3), and although the target throughput of 20–50 TPS is not yet reached, the stable TPS under increasing transaction volumes indicates linear scalability potential (DO4), suggesting realistic long-term fulfillment of both objectives. At the same time, achieving this throughput range in a privacy-preserving architecture critically depends on the efficiency of ZKP generation, which can become a computational bottleneck as transaction complexity increases (DO5). Because proof construction conditions how quickly transaction batches can be processed and finalized across layers, the scalability–privacy trade-off must be actively managed through use-case-specific circuit design that balances confidentiality requirements against computational overhead. In this regard, Aztec's PLONK-based circuit architecture and ongoing optimization of proving systems already aim to improve proof efficiency and support scalable private execution. Continued advances in proving performance and batching strategies remain essential to ensure that privacy-related computation does not limit effective TPS as procurement networks scale.

Beyond TPS, **latency** is critical for practical viability, as confirmation times determine usability and affect the feasibility of DO1, DO2, DO3, and DO5. In the proposed rollup-based architecture, latency must be distinguished into soft finality, which denotes the point at which a transaction is accepted by the sequencer and becomes visible to the application, and hard finality, which denotes the point at which it is cryptographically finalized on the underlying consensus layer [98]. Soft finality mainly impacts user experience and operational integration and thus relates to DO1 and DO2, whereas hard finality ensures trust, security, and auditability and thus relates to DO3 and DO5. For interorganizational applications such as procurement, soft-finality latency should remain within a few to several tens of seconds to preserve workflow continuity. Concretely, interactive steps such as negotiation actions should typically confirm within a few seconds and well below 10 s, while operational workflow steps such as fulfillment updates can tolerate latencies in the range of 10 to 60 s [99]. For settlement-related steps, soft-finality visibility within 1 to 2 min remains acceptable from a business-process perspective, as these steps are not part of a tight interaction loop. For hard finality, the current blockchain and enterprise practices expect settlement times of several minutes to at most a few tens of minutes [100]. In this context, finality within 30 min can be considered very good, while up to 60 min remains acceptable for accounting, auditing, and legally relevant settlement [101].

In our proposed design, Aztec achieves soft finality upon block inclusion and hard finality once the epoch is posted to L1 and the validity proof is verified. Since the current prototype operates in a simplified local setup, we refer to the public Ignition testnet for indicative latency expectations. Although the testnet currently produces epochs mainly containing empty blocks, Aztec's roadmap targets soft finality latencies of approximately 36–72 s [102], which would already suffice for fulfillment and settlement but not yet for sub-10-second interactive negotiation. The protocol's long-term goal of 4–12 s block times would achieve this [102,103], thus coming close to the soft-finality targets derived above and being critical for DO1 and DO2. For hard finality, current specifications and testnet benchmarks indicate epoch proving times of roughly 30–60 min [97], which aligns well with market expectations for cryptographic settlement and supports DO3 and DO5. Early audits and public protocol discussions suggest that this performance trajectory is technically plausible [104], particularly given the planned UltraHonk prover and decentralized sequencer architecture, and thus appears achievable in production.

Finally, beyond technical viability, the **economic sustainability** of the ecosystem must be assessed, since only a financially viable system can achieve long-term adoption. In the reference network, each invoice triggers about 7 verifiable events, totaling roughly 15,400 events per day. Practitioner reports indicate that manual invoice processing costs about $9–15 per invoice, while automation reduces this to roughly $3–5, implying savings of about $6–10 per invoice [105,106]. In a rollup-based architecture such as Aztec, blockchain-related costs arise from L2 transaction execution, sequencing, proof generation, and the amortized cost of posting and verifying compressed state updates on L1, where the total amortized cost per verifiable event refers to the combined cost across both layers. To derive an economically grounded threshold, we make the conservative design assumption that at most 10% of the achieved per-invoice savings should be spent on blockchain-related infrastructure, so that most of the benefit remains as net value after integration and operational overhead [107]. This yields an allowable cost budget of approximately $0.60–1.00 per invoice ($0.1 \times$ $6–10), which corresponds to about $0.09–0.14 per event ($0.60–1.00/7). This range defines a clear economic feasibility threshold that should not be exceeded to preserve the economic rationale for adoption.

Although no gas fees or payment mechanisms are currently implemented in either the prototype or the testnet, available documentation and analyses by the Aztec team suggest that transaction costs could be substantially lower than on Ethereum L1. According to Aztec's fee

Table 4
Quantitative viability assessment results.

Dimension	Metric	Derived Target Requirement	Prototype (local, Docker-based)	Aztec Testnet Documentation	Assessment
Scalability	Throughput (Transactions per Second)	20–50 TPS	0.05–0.07 TPS	Testnet currently produces empty blocks	Architecture shows linear scaling potential
Latency	Soft finality	Interactive: < 10 s Operational: < 60 s Settlement: 1–2 min	Not meaningfully measurable in the current setup	Current: 36–72 s Long-term goal: 4–12 s	Sufficient for operations today; interactive use depends on roadmap
	Hard finality	30–60 min	No reliable data (1 tx = 1 block in the artifact)	~ 30–60 min epoch proving time	Aligned with enterprise settlement expectations
Economic sustainability	Cost per event	≤ $0.09–0.14	No fee model implemented	~ $0.04 per event in an indicative example	Promising, but highly dependent on gas prices and batching

design, protocol fees are intended to account for both L1 execution and data availability costs as well as L2 sequencing and proof generation through a unified mechanism. In a published cost analysis of Aztec, the protocol team reports an example in which a privacy-preserving interaction required roughly 20× less L1 gas than the equivalent direct Ethereum transaction, illustrating the potential of the architecture to significantly amortize L1 costs through batching, and proof aggregation [108]. Assuming the gas usage reported in this example of about $24,672$ gas, together with a representative gas price in January 2026 of about 0.5 gwei and an ETH price of approximately $2900 [109], this would correspond to a cost of roughly $0.04 per event ($24,672 \times 0.5 \times 10^{-9} \times 2900$). While absolute costs depend on prevailing gas prices, this calculation illustrates that, under conditions such as sufficient batching, per-event costs could plausibly fall below the upper bound of the previously defined economic target range of approximately $0.09-0.14. However, since the current testnet does not yet operate at scale, these cost estimates remain indicative and depend on future network conditions and protocol evolution.

In summary, the quantitative evaluation indicates that the proposed architecture is technically and economically plausible with respect to throughput, latency, and cost under realistic assumptions, as summarized in Table 4. However, these results are based on estimates and early measurements and therefore establish directional feasibility rather than definitive performance guarantees, requiring further empirical validation in future evaluations. The following section, therefore, complements this analysis with a qualitative, expert-based evaluation against the predefined DOs.

7.2. Qualitative expert evaluation

To illustrate how verbatim expert feedback was systematically translated into assessments of DOs fulfillment, we provide an example for DO3, which concerns strengthening trust in procurement's dispute-resolution processes through permissionless infrastructure. After the demonstration of the artifact, one expert stated:*"I think it is fulfilled overall—at least in principle. However, that ultimately depends on the concrete implementation"*. (E10). This statement was coded as perceived fulfillment of permissionless trust and implementation-dependent trust, reflecting agreement with the design intent while emphasizing conditionality, resembling a half-filled circle within Table 5. Similarly, another expert highlighted the role of the permissionless base layer while identifying a specific architectural concern: *"The system is clearly permissionless at its core, anchored in an openly accessible Layer 1. The main aspect to examine more closely is whether the sequencer could become a single point of failure"*. (E13). This statement was coded as a permissionless trust anchor and as residual centralization risk at L2, and thus as partially achieved, because the expert feedback combines unconditional agreement with recurring, bounded concerns related to specific implementation aspects. In contrast, a third expert expressed unqualified agreement: *"This is fulfilled—completely and without reservation"*. (E2). This statement was coded as unconditional trust fulfillment, indicating that the artifact was perceived as fully satisfying the intent of DO3 without requiring additional qualifications. During axial consolidation, these codes were mapped to the higher-level evaluation category being trust enhancement through permissionless infrastructure. This assessment reflects qualified aggregation rather than subjective averaging and illustrates how full and conditional expert affirmations jointly inform the final evaluation outcome reported in Table 5.

In the context of DO1, experts agreed that the objective was addressed adequately but also highlighted some possible additional considerations. First, they noted that combining Aztec's APIs with the free development and deployment of smart contracts is a viable approach to fostering standardization (E2, E7, E8). According to E16, *"these interfaces are publicly accessible, allowing anyone to build on them and connect existing systems without prior permission or participation in a consortium"*. E5 further interpreted this as *"a common yet highly accessible implementation path, similar to how industry leaders define APIs that others adopt over time, with the difference that this path is more openly accessible"*. Furthermore, compared to closed or consortium-based solutions, the use of a public blockchain was seen as advantageous, as it enables direct and open ERP integration without onboarding constraints. E2, E7, and E16 further emphasized Ethereum's high interoperability as a key enabler for broad adoption and cross-organizational compatibility. E3 and E4, however, expressed concerns about Aztec's limited EVM compatibility, noting that this is often a critical consideration in the industry and could hinder integration, requiring further analyzes of possible consequences. E11 and E14 further observed that while the artifact performs well in standardized payment and procurement scenarios, it might show limitations in handling highly customized or dynamic information flows due to the static nature of smart contracts. E13 raised similar concerns but added that *"relying on a public and shared infrastructure may, over time, enable greater flexibility and modular extensibility"*. Overall, the experts evaluated the artifact positively. Several remarked that in certain customized use cases or integrations, the achievement of DO1 might require reevaluation and potential adjustments to the artifact. However, despite these caveats, there was a clear consensus that the artifact adequately addresses DO1.

Moreover, almost all experts agreed that the artifact contributes to the fulfillment of DO2, achieving stable and AML-compliant payments, by leveraging stablecoins governed by a MiCAR regulatory framework (E11, E12, E14, E14, E16). E4 further valued the artifact's construction, noting that *"for blockchain solutions to be effective, regulatory compliant payments must be integrated into workflows to prevent the emergence of separate payment silos, with stablecoins being a suitable means to accomplish*

Table 5
Expert evaluation of the artifact's alignment with the DOs.

	DO1	DO2	DO3	DO4	DO5
E1	●	◑	●	◑	◑
E2	◑	◑	◑	●	◑
E3	●	◑	◑	◑	◑
E4	●	●	◑	●	◑
E5	◑	◑	●	●	◑
E6	●	◑	●	●	◑
E7	◑	◑	◑	◑	●
E8	◑	●	◑	◑	●
E9	◑	●	●	◑	●
E10	●	◑	●	●	◑
E11	●	◑	◑	◑	◑
E12	◑	◑	●	●	◑
E13	●	●	●	●	◑
E14	●	◑	●	●	◑
E15	◑	◑	◑	◑	◑
E16	◑	◑	◑	●	◑

◑: fully achieved; ●: partially achieved;

this". However, while confident about AML compliance, several experts were uncertain about the specific role of CFT within MiCAR framework and the concrete requirements. Nevertheless, as highlighted by E16, *"regardless of the regulatory framework, whitelisting function is essential for high-value B2B payments"*. The main concern raised by some experts focused on potential errors and the irreversibility of blockchain-based transactions (E4, E8, E9). E13 added that whitelisting alone might not offer sufficient security, questioning what would happen in the case of a failure or flawed smart contracts. Consequently, although the design objective was considered fulfilled, several experts emphasized that regulatory safeguards and the potential consequences of transaction errors or smart contract failures must be more thoroughly evaluated and addressed in complex real-world settings.

Experts' opinions on the achievement of DO3 were mixed, yet many indicated that it holds promising potential if the underlying decentralization mechanisms function as intended. Several experts agreed that using a public-permissionless network, is a promising approach to strengthening trust in the dispute resolution process and the underlying infrastructure (E2, E3, E7, E15, E16). E15 also highlighted *"that a decentralized network can offer greater security compared to centralized alternatives"*. At the same time, E5 and E14 noted that FI could represent a single point of failure, as it controls access and is responsible for token minting. E10 and E13 further raised questions about the decentralization of the sequencer network, pointing out that *"while Aztec states the sequencer will be decentralized, this has yet to be demonstrated"* (E10). Even in the event of decentralization, the selection of sequencer could still affect the core security premise of a permissionless network (E10). Given the current developmental stage of the Aztec network, several experts considered it difficult to fully evaluate DO3. Nonetheless, they emphasized that if Aztec's decentralization efforts succeed as intended, the artifact represents a robust solution for balancing permissionless infrastructure with competing goals such as privacy and regulatory compliance.

The interview-based evaluation of the DO4 design revealed broad agreement among experts that scalability had been achieved. Experts (E7, E8, E9, E15) considered L2 solutions to be the most suitable approaches for achieving scalability today. The fact that the solution is built on Ethereum, which employs Proof-of-Stake and actively supports L2 technologies to enhance scalability, also received positive feedback (E1, E11). Furthermore, E1 highlighted that the *"PLONK protocol is currently one of the most efficient ZKPs systems"*, thereby mitigating common concerns about the computational overhead of proof generation raised by E12. While the underlying L2infrastructure was not questioned as a viable means to achieve scalability, some concerns were raised about the implementation layered on top. Experts E6, E13, and E14 observed that although the solution enhances scalability across several processes, it is less effective during the negotiation process. According to E13, *"even though bundling improves efficiency, continuously posting parts of the negotiation on-chain, only to have them stored by nodes that cannot access or interpret the data, introduces unnecessary inefficiencies"*. To address this, E6 and E13 suggested that smart contracts store only the final outcome of negotiations on-chain, noting that this can reduce auditability required in the DO4. Consequently, while the artifact's infrastructure was evaluated as strongly scalability-enhancing, experts emphasized that, depending on organizational trade-offs between auditability and scalability, the negotiation phase could be adapted in certain use cases to further enhance scalability.

The next evaluation step of DO5 yield clear and consistently positive results. All experts agreed that the most relevant information could be effectively concealed while using a public permissionless blockchain. Feedback was generally concise, with several experts confirming that we *"were able to fully achieve the DO5"* (E16). The only concern raised was related to future security. Once information is stored on-chain, even in encrypted and bundled form, it remains immutable (E7, E8, E9). As E8 noted, *"future technological advancements, such as quantum computing, might eventually make it possible to decrypt sensitive data"*. For this reason, applying ZKPs for privacy preservation requires additional evaluation of post-quantum resistance capabilities, which are typically not provided by the PLONK protocol but rather by alternatives such as STARK. Consequently, all experts agreed that the artifact successfully demonstrated a valid approach to fulfilling DO5. Nevertheless, some experts emphasized that long-term security considerations should form an integral part of protocol selection in the future.

8. Design principles and discussion

8.1. Design principles

The prototype developed in this study was evaluated quantitatively and qualitatively against five DOs that were themselves grounded in a detailed analysis of the information exchange challenges of procurement processes. The positive evaluation indicates that a public-permissionless base layer, combined with an Aztec privacy roll-up, can satisfy the identified objectives and thereby reduce media discontinuities, process inefficiencies, and trust deficits between procurement

Table 6
Design Principle 1.

Design Principle	**1. Principle of balancing the blockchain trilemma with a public Layer 1 as a trust anchor, a Layer 2 for scalability, and decentralized communication between layers.**	
Aim, implementer, and user	For developers and researchers of blockchain-based solutions (implementers), addressing business use cases such as procurement (users), aiming to leverage their solution by achieving a sustainable balance between scalability, security, and decentralization (aim).	
Context	In inter-organizational settings with trust, discontinuity, and inefficiency challenges.	
Mechanism	**1.1 Build the design on a public permissionless Layer 1 blockchain, complemented by a Layer 2 protocol.**	**1.2 Ensure that the Layer 2 design integrates a decentralized communication infrastructure.**
Rationale	Because it improves scalability while maintaining decentralization and security of the blockchain-based information system.	Because it ensures decentralization across layers, instead of shifting centralization risks from the base Layer 1 to the scaling Layer 2, thereby retaining the security guarantees of Layer 1.

partners. These results complement a growing body of procurement-oriented blockchain research. Earlier work has offered important contributions: some studies introduce prototypes that tackle discrete problems, such as media discontinuity [70,84], while others extract design guidance that encourages the use of user-friendly interfaces, increased process automation, and a careful assessment of the contextual factors that determine whether blockchain is an appropriate solution [74,79]. Our work builds on these contributions but shifts the analytical focus from the domains' process logic to the infrastructural design decisions. The findings derived from the prototype, thus, do not add new in-depth domain knowledge about procurement processes themselves but instead explain how specific architectural choices allow an information system based on blockchain for application in inter-organizational information exchange relationships to balance decentralization, scalability, security, and how this balance can be maintained when privacy requirements become stringent. The aim of this discussion is therefore to provide a solid initial foundation that can be empirically tested, critically adapted, and systematically extended as business contexts and technologies evolve. Building on this perspective, we propose DPs formulated in line with the approach suggested by Gregor et al. (2020), which focuses on the blockchain inherent infrastructural trade-offs and how they can be managed [39]. The developed principles aim to guide researchers and developers of blockchain-based business solutions, such as procurement, in enhancing the applicability of their prototypes with the meta goal of balancing the blockchain trilemma and maintaining this balance while adding privacy as an additional prerequisite. In this way, our work extends existing research by showing how infrastructural design can reinforce established blockchain design approaches and create additional value.

Design Principle 1 (DP1) (Table 6) addresses a central gap in blockchain-based procurement research concerning how decentralization, scalability, and security can be jointly balanced in inter-organizational settings without defaulting to permissioned or consortium-based architectures. While prior studies often emphasize individual dimensions of the blockchain trilemma, our design shows how these competing requirements can be reconciled through infrastructural design choices in permissionless environments. Thus, DP1 advises developers and researchers on how to balance the blockchain trilemma when designing blockchain-based information systems for procurement and comparable business processes characterized by trust deficits, media discontinuities, and operational inefficiencies. As illustrated by the artifact, a public-permissionless infrastructure fosters standardization (DO1) and strengthens trust (DO3), while L2 rollups mitigate scalability challenges (DO4), thus contributing to efficient payments (DO2). Building on this, the design principle distinguishes between two mechanisms: one aimed at enhancing scalability through L2 protocols built on a public-permissionless L1 infrastructure, and another focused on preserving decentralization and the security achieved by public-permissionless L1 by means of a decentralized communication networks among layers, such as the sequencer network in the prototype.

As outlined in DO5, privacy often plays a crucial role in business cases such as procurement. However, incorporating privacy can destabilize the achieved balance of the trilemma elements, as frequently highlighted during the expert interviews. Consequently, **Design Principle 2 (DP2)** (Table 7) addresses a complementary gap concerning the integration of privacy into permissionless blockchain architectures. Existing procurement research often treats privacy as a justification for restricting participation or centralizing control, providing limited guidance on how to incorporate privacy-enhancing technologies without compromising decentralization, scalability, or security. DP2 specifies how cryptographic mechanisms and minimized control structures preserve the extended blockchain trilemma under stringent confidentiality requirements, guiding developers and researchers of blockchain-based procurement information systems in integrating privacy without undermining this balance. The principle introduces three mechanisms. The first two concern the selection of ZKPs protocols to safeguard scalability and security. While ZKPs provide an effective means of protecting sensitive data, as demonstrated by the artifact, proof generation can be energy-intensive, and generated proofs may create future security risks. Furthermore, proving and verification overhead may influence achievable throughput and latency under high transaction volumes, as elaborated in the quantitative evaluation section. Developers must therefore choose efficient and secure protocols that prevent information leakage to the base and communication layers.

Yet, a trade-off remains: efficient protocols such as PLONK lack quantum resistance, whereas quantum-resistant protocols such as STARK are comparatively slower [110]. In the developed artifact, this trade-off primarily affects the inter-layer communication, which relies on cryptographic proofs, since private execution logs are only stored temporarily on L1. While this design limits the risk of compromising the entire transaction history, it does not eliminate the long-term risk associated with transmitted data. Consequently, depending on the sensitivity of the use case, it is important to consider long-term mitigation strategies should quantum computers become practical. Such strategies may include selectively using STARK-based proofs for highly sensitive data or adopting hybrid constructions. Since Noir is backend-agnostic, developers can switch between proof systems where necessary [92]. Finally, combining privacy with regulatory requirements may threaten decentralization by introducing control mechanisms into the smart contract ecosystem. To safeguard decentralization, such structures should be kept to a minimum.

Table 7
Design Principle 2.

Design Principle	**2. Principle of ensuring privacy while maintaining a balanced blockchain trilemma through efficient and resilient cryptography and minimized control mechanism.**		
Aim, implementer, and user	For developers and researchers of blockchain-based solutions (implementers), addressing use cases such as procurement (users), aiming to leverage their solution by integrate privacy while maintaining the equilibrium of the scalability, security, and decentralization trilemma (aim).		
Context	In inter-organizational settings with trust, discontinuity, and inefficiency challenges, where privacy is additionally required.		
Mechanism	**2.1 Ensure that the chosen Layer 2 solution applies efficient ZKPs, with proof generation occurring in a secure off-chain environment.**	**2.2 Ensure that the chosen Layer 2 solution relies on ZKP protocols that provide long-term cryptographic resilience.**	**2.3 Minimize reliance on centralized control structures within the smart contract ecosystem.**
Rationale	Because it safeguards privacy without compromising scalability, by ensuring confidential data handling in a secure environment, such as a Trusted Execution Environment, while enabling efficient proof generation.	Because it safeguards privacy without undermining security, ensuring that tamper-resistant and encrypted data on the blockchain remain resilient against potential adversaries, including quantum attacks.	Because it safeguards privacy without undermining decentralization, avoiding permissioned settings insofar as this remains consistent with prevailing regulatory obligations.

8.2. Discussion

This study contributes to blockchain scholarship on procurement at three inter-linked levels: an empirically validated problem statement, a set of literature- and expert-derived DOs, and two prescriptive DPs distilled from an evaluated artifact. Taken together, these elements demonstrate that a layered architecture, public-permissionless L1 combined with a privacy-preserving roll-up, can reconcile decentralization, scalability, security, and privacy, thus leveraging routine procure-to-pay workflows.

First, the problem statement itself advances prior research. Earlier studies often examined procurement inefficiencies and blockchain design tensions in isolation, leaving their interrelation open for future investigation [27,43,84]. Our SLR integrates these perspectives by outlining key procurement challenges that blockchain can address, analyzing the types of blockchains applied in previous studies, and highlighting infrastructural gaps that remain underexplored. The results therefore show that frequently used permissioned infrastructures are often insufficient to resolve scalability issues or to provide the full decentralization-based security, thus limiting the potential advantages of developed solutions [64,65].

Second, the five DOs translate this problem framing into technically actionable requirements, demonstrating how procurement challenges can be addressed through blockchain technology, while also benefiting from the consideration of balanced trilemma elements. Developed and iteratively refined with practitioners, the objectives reflect real-world priorities, such as regulatory-compliant token payments and standardization for improved interoperability, that purely conceptual models often overlook [71,75]. As intermediate constructs, the evaluated DOs ground the artifact's design and provide a reusable template for future studies.

Third, the artifact's evaluation informs two DPs that speak directly to ongoing debates in blockchain research. Both principles challenge the dominant view that confidentiality and throughput can be achieved mainly on permissioned or consortium chains. By anchoring the critical state on a public base layer while delegating volume to a decentralized roll-up, DP1 provides performance without recentralizing control, thereby complementing the long-standing "public-versus-private" dichotomy reported in enterprise blockchains solutions [23,28]. DP2 embeds privacy within this architecture through efficient ZKPs executed in secure off-chain environments and through a minimized reliance on central control structures. This approach advances the literature from a "privacy-or-performance" trade-off to a design that accommodates both, empirically substantiating conceptual calls for privacy-by-design blockchains [110].

Especially for practitioners, the results, thus, translate into several interrelated recommendations. Recording procurement events on a public chain secures non-repudiation and future-proof interoperability, while processing high-volume or sensitive data in a decentralized roll-up maintains throughput and confidentiality. Sequencer roles within the roll-up should be shared or rotated among stakeholders to avoid new single points of failure, and governance rules for validator admission and dispute resolution need to be settled early. When selecting ZKPs, organizations must balance proof cost against long-term cryptographic resilience. PLONK variants may suit high-throughput networks, whereas STARK-based schemes offer stronger post-quantum security at higher computational expense.

Beyond the EU regulatory context exemplified by MiCAR and GDPR, real-world deployments may face stricter data localization regimes, enhanced supervisory oversight, or jurisdiction-specific reporting obligations. The modular separation between the permissionless trust anchor at the base layer and jurisdiction-specific compliance components at the application or roll-up level enables configuration-level adaptations without altering the core architecture. Sensitive procurement data can remain within jurisdictionally bounded environments, while only cryptographic commitments are anchored to L1, thereby supporting compliance with data sovereignty requirements without sacrificing cross-border verifiability. Governance parameters, such as participation criteria, regulatory visibility mechanisms, or sequencing arrangements, can be adjusted to reflect varying degrees of supervisory intervention. This architectural modularity enhances the design's global applicability across heterogeneous regulatory environments while preserving decentralization, auditability, and scalability.

Finally, the study not only advises practitioners but also contributes to theory by extending the blockchain trilemma to include privacy as a fourth axis and showing how all dimensions are interrelated. It further broadens blockchain privacy research, which has so far focused mainly on payments or single-party data storage, by validating a ZKPs-enabled roll-up in a multi-party procurement scenario. As such, this paper exemplifies how design-science research can advance both theoretical understanding and practical design of blockchain infrastructures for business cases such as procurement.

9. Conclusion

This paper set out to demonstrate how procurement, one of the most data and coordination-intensive functions in supply chains, can benefit from a blockchain-based information system that reconciles decentralization, scalability, security, and privacy. Guided by the DSR paradigm [36], we first synthesized an empirically validated problem

statement from the literature and expert interviews, then translated it into five DOs, and finally instantiated and evaluated a prototype that anchors trust on a public permissionless blockchain while delegating high-volume and confidential processing to a privacy-preserving roll-up. The quantitative evaluation was based on estimated requirements for a mid-sized supply chain company and examined scalability, latency, and cost efficiency to assess the feasibility of real-world deployment. While the results indicate promising potential, they require further validation as the underlying Aztec infrastructure layer matures. Given the assumption-based nature of this evaluation, we complemented it with a qualitative study based on interviews with domain experts to triangulate the findings and assess their practical validity. The evaluation shows that the prototype addresses all DO in a realistic procure-to-pay flow, while also revealing that their realization is subject to specific architectural, organizational, and governance conditions, and that some objectives involve design trade-offs that may be optimized for specific organizational contexts. In particular, expert feedback suggests that achieving interoperability and standardization (DO1) depends on the maturity of integration standards and ecosystem adoption, while enhancing trust through permissionless infrastructure (DO3) is contingent upon the effective decentralization of key components, such as sequencing and governance mechanisms. Reflecting on these build–evaluate cycles, we distilled two prescriptive design principles that make such conditions explicit and actionable, thereby providing guidance for designing blockchain-based procurement systems that balance the extended blockchain trilemma in practice.

However, like every research effort, the present work is subject to limitations that point to fruitful avenues for future inquiry. A first limitation concerns technological maturity. Aztec is still under active development, with several roadmap features, including batching strategies for large proof sets and formal verification of the new Noir language, being unavailable during the preparation of our study. Consequently, we were unable to benchmark the throughput of proof generation under peak loads or test the security guarantees against sophisticated adversaries. This leaves scope for future research into the artifact's production-level robustness. In addition, once more mature implementations become available, future work should as well conduct systematic quantitative and qualitative comparisons with alternative privacy-focused scaling solutions to assess their relative strengths and the extent to which the prototype would need to be adapted to other suitable platforms.

A second limitation lies in regulatory scope. Because the roll-up is settled on a public blockchain, any real-world deployment must comply with financial and data protection laws that vary by jurisdiction. Our prototype incorporated only the MiCAR-compatible stablecoin layer that the participating experts deemed most urgent. Broader compliance requirements, such as stricter data localization regimes, enhanced supervisory oversight, or jurisdiction-specific reporting obligations, may necessitate configuration-level adaptations, including selective re-encryption, data retention schedules, or modified governance parameters at the application or roll-up layer. Although we discuss regulatory extensibility beyond MiCAR at a conceptual level (cf. Section 6.2 and Section 8), a comprehensive empirical evaluation of jurisdiction-specific legal requirements (e.g., GDPR enforcement practices or non-EU procurement regulations) is beyond the scope of this study and represents an important avenue for future research on real-world deployment.

Third, the design-science methodology introduces its own constraints: the insights stem from one instantiated artifact and expert interviews for evaluation. Beyond technical and regulatory considerations, real-world deployment may also face practical challenges, including adoption and integration costs, user onboarding and change management efforts, and interoperability with heterogeneous legacy (procurement) systems. These aspects were outside the scope of this study but are critical for successful industrial adoption. While the DSR approach yields depth and relevance, it cannot by itself quantify efficiency gains or capture the full diversity of organizational contexts.

These limitations suggest several avenues for future research. Large-scale field trials should evaluate the architecture's performance and governance under real operating conditions. Beyond procurement, future research could test the scalability and transferability of the proposed design principles in other data-intensive, inter-organizational domains. Promising contexts include logistics and supply chain coordination, where high transaction volumes and multi-party visibility are critical, as well as healthcare, where stringent privacy requirements and regulatory oversight coexist with the need for shared, tamper-evident records. Evaluating the principles in such settings would allow researchers to assess their robustness under different workload characteristics, regulatory regimes, and data-sensitivity profiles. In addition, extending the artifact with automated dispute resolution, oracle-based event attestation, and multi-tier supplier visibility would test the design principles in more complex supply-chain ecosystems. Second, rigorous cost–benefit studies are needed: Quantitative surveys, process-mining analyses, and comparative assessments against traditional EDI or permissioned-chain systems can quantify efficiency gains, implementation costs, and ongoing operating expenses, thus informing managerial decision making. Comparative studies could also benchmark alternative privacy-enhancing tools against ZKPs roll-ups to refine the choice of cryptographic primitives. Despite these open questions, this study provides an empirically validated architecture and DPs that offer scholars and practitioners a solid foundation for deploying secure, scalable, and privacy-preserving blockchain systems in procurement.

CRediT authorship contribution statement

Valeriya Arnold: Writing – review & editing, Writing – original draft, Visualization, Validation, Software, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Tobias Guggenberger:** Writing – review & editing, Supervision, Project administration, Investigation, Conceptualization. **Jan Stramm:** Writing – review & editing, Writing – original draft, Visualization, Project administration, Methodology, Formal analysis, Data curation, Conceptualization. **Nils Urbach:** Writing – review & editing, Supervision, Project administration, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

References

[1] R. Monczka, R. Handfield, L. Giunipero, Purchasing and Supply Chain Management, sixth ed., Cengage Learning, Florence, AL, 2014.
[2] M. Kibria, B. George, Blockchain applications to mitigate the effects of supply chain disruptions, Am. Conf. Inf. Syst. 2022 (2022) 1–5.
[3] S. Konecka, Z. Bentyn, Cyberattacks as threats in supply chains, Eur. Res. Stud. J. XXVII (3) (2024) 778–796.
[4] D. Ozdemir, M. Sharma, A. Dhir, T. Daim, Supply chain resilience during the COVID-19 pandemic, Technol. Soc. 68 (2022) 1–10.
[5] J.T. Mentzer, W. DeWitt, J.S. Keebler, S. Min, N.W. Nix, C.D. Smith, Z.G. Zacharia, Defining supply chain management, J. Bus. Logist. 22 (2) (2001) 1–25, http://dx.doi.org/10.1002/j.2158-1592.2001.tb00001.x.
[6] M.C. Cooper, D.M. Lambert, J.D. Pagh, Supply chain management: More than a new name for logistics, Int. J. Logist. Manag. 8 (1) (1997) 1–14, http://dx.doi.org/10.1108/09574099710805556.
[7] M. Shi, W. Yu, Supply chain management and financial performance: literature review and future directions, Int. J. Oper. Prod. Manage. 33 (10) (2013) 1283–1317, http://dx.doi.org/10.1108/IJOPM-03-2012-0112.
[8] R.D. Nelson, P.E. Moody, J. Stegner, The Purchasing Machine: how the Top Ten Companies Use Best Practices to Manage Their Supply Chains, Free Press, New York, NY, 2001.

[9] A. Gunasekaran, E. Ngai, Information systems in supply chain integration and management, European J. Oper. Res. 159 (2) (2004) 269–295, https://doi.org/10.1016/j.ejor.2003.08.016.

[10] D. Ivanov, A. Dolgui, B. Sokolov, Cloud supply chain: Integrating industry 4.0 and digital platforms in the "supply chain-as-a-service", Transp. Res. Part E: Logist. Transp. Rev. 160 (2022) 1–11, http://dx.doi.org/10.1016/j.tre.2022.102676.

[11] S. Croom, A. Brandon-Jones, Impact of e-procurement: Experiences from implementation in the UK public sector, J. Purch. Supply Manag. 13 (4) (2007) 294–303, http://dx.doi.org/10.1016/j.pursup.2007.09.015.

[12] V.H. Villena, The missing link? The strategic role of procurement in building sustainable supply networks, Prod. Oper. Manage. 28 (5) (2019) 1149–1172, http://dx.doi.org/10.1111/poms.12980.

[13] R. Alt, M. Gräser, Distributed ledger technology, Electron. Mark. 35 (1) (2025) http://dx.doi.org/10.1007/s12525-025-00784-w.

[14] A. Atadoga, F. Osasona, O.O. Amoo, O.A. Farayola, B.S. Ayinla, T.O. Abrahams, The role of IT in enhancing supply chain resilience: a global review, Int. J. Manag. Entrep. Res. 6 (2) (2024) 336–351.

[15] R. Manzoor, B. Sahay, S.K. Singh, Blockchain technology in supply chain management: an organizational theoretic overview and research agenda, Ann. Oper. Res. (2022) 1–48, https://doi.org/10.1007/s10479-022-05069-5.

[16] C. Finke, M. Alfken, M. Schumann, Why do distributed ledger platforms fail? Analyzing the challenges of distributed ledger technologies in supply chain processes, Am. Conf. Inf. Syst. 2023 (2023) 1–20.

[17] J. Kolb, L. Becker, M. Fischer, A. Winkelmann, The role of blockchain in enterprise procurement, Hawaii Int. Conf. Syst. Sci. 2019 (2019) 1–10.

[18] J. Wijayasiri, Implementation of TradeLens in Sri Lanka using blockchain technology: a case study, UNNExT Work. Pap. Ser. No. 1 (2023) 1–19.

[19] H. Partz, Tencent-backed everledger collapses amid lack of funding: Report, 2023, https://cointelegraph.com/news/tencent-backed-everledger-collapses-amid-lack-of-funding-report (Accessed 14 June 2025).

[20] IBM, An IBM blockchain solution for the complete food system, 2024, https://www.ibm.com/docs/en/food-trust (Accessed 7 January 2025).

[21] De Beers Group, Blockchain technology and diamond traceability, 2024, https://www.debeersgroup.com/about-us/case-studies/2024/tracr (Accessed 13 January 2025).

[22] M. Kuciapski, F. Nazet, Blockchain technology perception regarding supporting the digital transformation of supply chain management, Am. Conf. Inf. Syst. 2023 (2023).

[23] D. Mishunin, Why blockchain falls short in supply chain management: Here's why, 2024, https://www.forbes.com/councils/forbestechcouncil/2024/01/08/why-blockchain-falls-short-in-supply-chain-management-heres-why/ (Accessed 8 January 2024).

[24] A. Hafid, A.S. Hafid, M. Samih, Scaling blockchains: A comprehensive survey, "IEEE Access" 8 (2020) 125244–125262.

[25] M. Principato, M. Babel, T. Guggenberger, J. Kropp, S. Mertel, Towards solving the blockchain trilemma: An exploration of zero-knowledge proofs, Forty-Forth Int. Conf. Inf. Syst. Hyderabad 2023 (2023) 1–18.

[26] M. Kim, Y.-W. Kim, Applications of blockchain for construction project procurement, Autom. Constr. 165 (2024) 1–22, https://doi.org/10.1016/j.autcon.2024.105550.

[27] H. Ritchi, A. Bandana, Z. Adrianto, A. Alfian, Permissioned blockchain for business process visibility: A case of expenditure cycle, Procedia Comput. Sci. 197 (2022) 336–343, https://doi.org/10.1016/j.procs.2021.12.148.

[28] T. Nodehi, A. Zutshi, A. Grilo, B. Rizvanovic, EBDF: The enterprise blockchain design framework and its application to an e-procurement ecosystem, Comput. Ind. Eng. 171 (2022) 1–20, https://doi.org/10.1016/j.cie.2022.108360.

[29] R. Yang, R. Wakefield, S. Lyu, S. Jayasuriya, F. Han, X. Yi, X. Yang, G. Amarasinghe, S. Chen, Public and private blockchain in construction business process and information integration, Autom. Constr. 118 (2020) 1–21, https://doi.org/10.1016/j.autcon.2020.103276.

[30] B.P.C. Teoh, Navigating the blockchain trilemma: a supply chain dilemma, Adv. Marit. Technol. Appl.: Pap. from ICMAT 2021 (2022) 291–300, https://doi.org/10.1007/978-3-030-89992-9_25.

[31] S. Goldwasser, S. Micali, C. Rackoff, The knowledge complexity of interactive proof-systems, in: Providing Sound Foundations for Cryptography: On the Work of Shafi Goldwasser and Silvio Micali, 2019, pp. 203–225, https://doi.org/10.1145/3335741.3335750.

[32] P. Singh, M. Masud, M.S. Hossain, A. Kaur, Blockchain and homomorphic encryption-based privacy-preserving data aggregation model in smart grid, Comput. Electr. Eng. 93 (2021) 1–11, https://doi.org/10.1016/j.compeleceng.2021.107209.

[33] L.T. Thibault, T. Sarry, A.S. Hafid, Blockchain scaling using rollups: A comprehensive survey, IEEE Access 10 (2022) 93039–93054, https://doi.org/10.1109/ACCESS.2022.3200051.

[34] C. James, M. Tudor, M. Cais, N. Shah, T. Gavin, Obscure whitepaper, 2024, https://whitepaper.ten.xyz/assets/images/obscuro-whitepaper-0-10-0.pdf (Accessed 14 December 2024).

[35] Aztec, Aztec overview, 2025, https://docs.aztec.network/nightly/developers/docs/concepts (Accessed 14 September 2025).

[36] K. Peffers, T. Tuunanen, M.A. Rothenberger, S. Chatterjee, A design science research methodology for information systems research, J. Manage. Inf. Syst. 24 (3) (2007) 45–77, https://doi.org/10.2753/MIS0742-1222240302.

[37] P.B. Kruchten, The 4+ 1 view model of architecture, IEEE Softw. 12 (6) (2002) 42–50, https://doi.org/10.1109/52.469759.

[38] J. Venable, J. Pries-Heje, R. Baskerville, FEDS: a framework for evaluation in design science research, Eur. J. Inf. Syst. 25 (1) (2016) 77–89, https://doi.org/10.1057/ejis.2014.36.

[39] S. Gregor, L. Chandra Kruse, S. Seidel, Research perspectives: the anatomy of a design principle, J. Assoc. Inf. Syst. 21 (6) (2020) 1622–1652, https://doi.org/10.17705/1jais.00649.

[40] A. Gunasekaran, E. Ngai, Data quality for data science, predictive analytics, and big data in supply chain management: An introduction to the problem and suggestions for research and applications, European J. Oper. Res. 154 (2) (2004) 269–295, http://dx.doi.org/10.1016/j.ijpe.2014.04.018.

[41] J. Hallikas, K. Puumalainen, T. Vesterinen, V.-M. Virolainen, Risk-based classification of supplier relationships, J. Purch. Supply Manag. 11 (2–3) (2005) 72–82, https://doi.org/10.1016/j.pursup.2005.10.005.

[42] G. Burtch, Y. Hong, S. Kumar, When does dispute resolution substitute for a reputation system? empirical evidence from a service procurement platform, Prod. Oper. Manage. 30 (6) (2021) 1565–1582, http://dx.doi.org/10.1111/poms.13341.

[43] M. Saygili, I.E. Mert, O.B. Tokdemir, A decentralized structure to reduce and resolve construction disputes in a hybrid blockchain network, Autom. Constr. 134 (2022) 1–17, https://doi.org/10.1016/j.autcon.2021.104056.

[44] B.-J. Butijn, D.A. Tamburri, W.-J.v.d. Heuvel, Blockchains: A systematic multivocal literature review, ACM Comput. Surv. 53 (3) (2020) 1–37, https://doi.org/10.1145/3369052.

[45] J. Abadi, M. Brunnermeier, Blockchain Economics, National Bureau of Economic Research, 2018, http://dx.doi.org/10.3386/w25407.

[46] B. Lashkari, P. Musilek, A comprehensive review of blockchain consensus mechanisms, IEEE Access 9 (2021) 43620–43652, http://dx.doi.org/10.1109/ACCESS.2021.3065880.

[47] E. Ellinger, R. Gregory, T. Mini, T. Widjaja, O. Henfridsson, Skin in the game: The transformational potential of decentralized autonomous organizations, MIS Q. 48 (1) (2024) 245–272, http://dx.doi.org/10.25300/MISQ/2023/17690.

[48] R. Beck, M. Avital, M. Rossi, J.B. Thatcher, Blockchain technology in business and information systems research, Bus. & Inf. Syst. Eng. 59 (6) (2017) 381–384, http://dx.doi.org/10.1007/s12599-017-0505-1.

[49] J. Lohmer, L. Petzok, R. Lasch, Governance design of blockchain consortia for efficient and transparent procurement and supply chain management, in: Supply Management Research, Springer Gabler, Wiesbaden, 2021, https://doi.org/10.1007/978-3-658-35449-7_6.

[50] L.F. Bossler, A. Buchwald, K. Spohrer, And no one gets the short end of the stick: A blockchain-based approach to solving the two-sided opportunism problem in interorganizational information sharing, Inf. Syst. Res. (2024) http://dx.doi.org/10.1287/isre.2022.0065.

[51] J.H. Yoon, I. Aurangzeb, S. McNamara, BIM- and blockchain-enabled automatic procurement system (BBAPS) removing relationship bias, Autom. Constr. 168 (2024) 1–19, https://doi.org/10.1016/j.autcon.2024.105779.

[52] S.N. Yutia, B. Rahardjo, Design of a blockchain-based e-Tendering system: A case study in LPSE, Int. Conf. ICT Smart Soc. (ICISS) (2019) 1–6, https://doi.org/10.1109/ICISS48059.2019.8969824.

[53] LBBW, LBBW and partners go live with marco polo trade finance network, 2021, https://www.lbbw.de/article/press-release/marco-polo-goes-live_ac27ox25iu_e.html (Accessed 7 July 2025).

[54] J. Lautenschlager, J. Stramm, T. Guggenberger, N. Urbach, Striking a balance: Designing a blockchain-based solution to navigate coopetition dynamics in supply chain management, "Electron. Mark." 35 (1) (2025) 1–24, https://doi.org/10.1007/s12525-025-00809-4.

[55] T. Guggenberger, J. Sedlmeir, G. Fridgen, A. Luckow, An in-depth investigation of the performance characteristics of hyperledger fabric, Comput. & Ind. Eng. 173 (2022) 1–20, http://dx.doi.org/10.1016/j.cie.2022.108716.

[56] R.G. Brown, J. Carlyle, I. Grigg, M. Hearn, Corda: An introduction, Work. Pap. (2016) http://rgdoi.net/10.13140/RG.2.2.30487.37284.

[57] S. Ding, H. Hu, F. Xu, Z. Chai, W. Wang, Blockchain-based security-minded information-sharing in precast construction supply chain management with scalability, efficiency and privacy improvements, Autom. Constr. 168 (2024) 1–23, https://doi.org/10.1016/j.autcon.2024.105698.

[58] J. Sedlmeir, J. Lautenschlager, G. Fridgen, N. Urbach, The transparency challenge of blockchain in organizations, Electron. Mark. 32 (3) (2022) 1779–1794, http://dx.doi.org/10.1007/s12525-022-00536-0.

[59] J. Zhou, Y. Feng, Z. Wang, D. Guo, Using secure multi-party computation to protect privacy on a permissioned blockchain, Sensors 21 (4) (2021) 1–17, https://doi.org/10.3390/s21041540.

[60] J. Webster, R.T. Watson, Analyzing the past to prepare for the future: Writing a literature review, MIS Q. 26 (2) (2002) xiii–xxiii, https://www.jstor.org/stable/4132319.

[61] Findora, Findora is a next-generation blockchain focused on protecting data privacy, 2024, https://mainnet.findora.org/ (Accessed 14 October 2024).

[62] Z.J. Williamson, The AZTEC protocol, 2018, https://github.com/AztecProtocol/aztec-v1/blob/master/AZTEC.pdf.

[63] J. Corbin, A. Strauss, Grounded theory methodology, Handb. Qual. Res. 17 (1) (1994) 273–285.

[64] T.I. Akaba, A. Norta, C. Udokwu, D. Draheim, A framework for the adoption of blockchain-based e-Procurement systems in the public sector, Responsible Des. Implement. Use Inf. Commun. Technol. (2020) 3–14, https://doi.org/10.1007/978-3-030-44999-5_1.

[65] P. Diadia, J.K. Tamgno, A.D. Kora, Implementing and evaluating a blockchain-based dematerialized public procurement system with HyperLedger fabric and OCDS, Fifth Int. Conf. Blockchain Comput. Appl. (BCCA) (2023) 136–141, https://doi.org/10.1109/BCCA58897.2023.10338876.

[66] T. Weingärtner, D. Batista, S. Köchli, G. Voutat, Prototyping a smart contract based public procurement to fight corruption, Computers 10 (7) (2021) https://doi.org/10.3390/computers10070085.

[67] S. Ahmadisheykhsarmast, R. Sonmez, A smart contract system for security of payment of construction contracts, Autom. Constr. 120 (2020) 1–14, https://doi.org/10.1016/j.autcon.2020.103401.

[68] M. Basheer, F. Elghaish, T. Brooks, F. Pour Rahimian, C. Park, Blockchain-based decentralised material management system for construction projects, J. Build. Eng. 82 (2024) 1–24, https://doi.org/10.1016/j.jobe.2023.108263.

[69] R.A. Ara, K. Paardenkooper, R. van Duin, A new blockchain system design to improve the supply chain of engineering, procurement and construction (EPC) companies - a case study in the oil and gas sector, J. Eng. Des. Technol. 20 (4) (2022) 887–913, https://doi.org/10.1108/JEDT-01-2021-0047.

[70] B.K. Rai, V. Dubey, K. Dubey, Blockchain based E-procurement system in healthcare, Health Serv. Outcomes Res. Methodol. 24.0 (4) (2024) 403–426, https://doi.org/10.1007/s10742-023-00321-2.

[71] T.S. Elbashbishy, G.G. Alil, I.H. El-adaway, Role of transactional blockchain in facilitating procurement in international construction projects, Proc. Can. Soc. Civ. Eng. Annu. Conf. 2021 (2022) 583–596, https://doi.org/10.1007/978-981-19-1029-6_44.

[72] G.M.P. Nechkoska, Conceptualisation of decentralized blockchain-based, open-source ERP marketplaces, in: Facilitation in Complexity, 2023, pp. 172–202, https://doi.org/10.1007/978-3-031-11065-8_7.

[73] L. Trautmann, R. Lasch, Blockchain-based smart contracts in procurement: A technology readiness level analysis, in: Einkauf Und Supply Chain Management, 2021, pp. 133–170, https://doi.org/10.1007/978-3-658-32895-5_6.

[74] J. Lautenschlager, J. Stramm, T. Guggenberger, A. Rösch, A. Schweizer, Overcoming the data transparency trade-off: Designing a blockchain-based delivery invoice system for the construction industry, Wirtsch. 2023 Proc. (2023) 377–392, https://doi.org/10.1007/978-3-031-80122-8_24.

[75] I.E. Inghirami, Accounting information systems: The scope of blockchain accounting, Digit. Bus. Transform. (2020) 107–120, https://doi.org/10.1007/978-3-030-47355-6_8.

[76] L.W.L. Xu, Blockchain-based smart contract for smart payment in construction: A focus on the payment freezing and disbursement cycle, Front. Eng. Manag. 9.0 (2) (2022) 177–195, https://doi.org/10.1007/s42524-021-0184-y.

[77] F. Elghaish, F. Pour Rahimian, M.R. Hosseini, D. Edwards, M. Shelbourn, Financial management of construction projects: Hyperledger fabric and chaincode solutions, Autom. Constr. 137 (2022) 1–14, https://doi.org/10.1016/j.autcon.2022.104185.

[78] S. Ahmadisheykhsarmast, S.G. Senji, R. Sonmez, Decentralized tendering of construction projects using blockchain-based smart contracts and storage systems, Autom. Constr. 151 (2023) 1–16, https://doi.org/10.1016/j.autcon.2023.104900.

[79] T. Henry, R. Beck, N. Laga, W. Gaaloul, S. Pan, Decentralized procurement mechanisms for efficient logistics services mapping - a design science research approach, Hawaii Int. Conf. Syst. Sci. (HICSS-55) (2022) 1–10.

[80] J.-L. Ferrer-Gomila, M. Francisca Hinarejos, A.-P. Isern-Deyà, A fair contract signing protocol with blockchain support, Electron. Commer. Res. Appl. 36 (2019) 1–13, https://doi.org/10.1016/j.elerap.2019.100869.

[81] H. Hamledari, M. Fischer, Construction payment automation using blockchain-enabled smart contracts and robotic reality capture technologies, Autom. Constr. 132 (2021) 1–14, https://doi.org/10.1016/j.autcon.2021.103926.

[82] J. Kolb, J. Hornung, F. Kraft, A. Winkelmann, Industrial application of blockchain technology – Erasing the weaknesses of vendor managed inventory, Eur. Conf. Inf. Syst. (ECIS) (2018) 1–17.

[83] Pham, Vu Hong Son, T.T. Vo, N.T.N. Dang, Applying blockchain technology in smart contracts for construction payment: a comprehensive solution for lumpsum contracts, Asian J. Civ. Eng. 25.0 (4) (2024) 3549–3564, https://doi.org/10.1007/s42107-024-00995-0.

[84] A. Banerjee, Blockchain technology: Supply chain insights from ERP, in: Advances in Computers : Blockchain Technology: Platforms, Tools and Use Cases, vol. 111, Elsevier, 2018, pp. 69–98, https://doi.org/10.1016/bs.adcom.2018.03.007.

[85] A. Faccia, P. Petratos, Blockchain, enterprise resource planning (ERP) and accounting information systems (AIS): Research on e-procurement and system integration, Appl. Sci. 11 (15) (2021) 1–17, https://doi.org/10.3390/app11156792.

[86] F. Sunmola, G.L. Lawrence, Key success factors for integration of blockchain and ERP systems: A systematic literature review, Procedia Comput. Sci. 232 (2024) 775–782, https://doi.org/10.1016/j.procs.2024.01.077.

[87] L. Li, J. Liu, P. Jia, SecTEP: Enabling secure tender evaluation with sealed prices and quality evaluation in procurement bidding systems over blockchain, Comput. Secur. 103 (2021) 1–15, https://doi.org/10.1016/j.cose.2021.102188.

[88] A. Gabizon, Z.J. Williamson, O. Ciobotaru, Plonk: Permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge, Cryptol. EPrint Arch. (2019) 1–34, https://eprint.iacr.org/2019/953.

[89] H. Liu, X. Luo, H. Liu, X. Xia, Merkle tree: A fundamental component of blockchains, Int. Conf. Electron. Inf. Eng. Comput. Sci. (EIECS) (2021) 556–561, https://doi.org/10.1109/EIECS53707.2021.9588047.

[90] Aztec, L1-l2 communication (portals), 2025, https://docs.aztec.network/aztec/concepts/communication/portals (Accessed 10 January 2025).

[91] Q. Wang, S. Chen, Account abstraction, analysed, IEEE Int. Conf. Blockchain (2023) 323–331, https://doi.org/10.1109/Blockchain60715.2023.00057.

[92] Noir, Noir language, 2025, https://noir-lang.org/docs/ (Accessed 10 January 2025).

[93] F. Annunziata, An overview of the markets in crypto-assets regulation (MiCAR), Eur. Bank. Inst. Work. Pap. Ser. No. 158 (2023) 1–72, http://dx.doi.org/10.2139/ssrn.4660379.

[94] European Commission, Evaluation of directive 2014/55/EU on electronic invoicing in public procurement, 2024, https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX%3A52024SC0039&utm (Accessed 10 January 2026).

[95] Eurostat, Structural business statistics overview, 2024, https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Structural_business_statistics_overview(Accessed 10 January 2026).

[96] L.Q. Torres, D. Colish, SRE best practices for capacity management, Proc. USENIX PATRONS (2020) 1–49.

[97] Aztec, Enter the adversarial testnet arena, 2026, https://testnet.aztec.network/ (Accessed 15 January 2026).

[98] K. Gogol, M. Schneider, C. Tessone, First-spammed, first-served: MEV extraction on fast-finality blockchains, 2025, arXiv preprint arXiv:2506.01462.

[99] J. Nielsen, Response times: The 3 important limits, 2026, https://www.nngroup.com/articles/response-times-3-important-limits/?utm (Accessed 14 January 2026).

[100] K. Gogol, S. Gurgul, F.N. Siddiqui, D. Branes, C. Tessone, Scaling DeFi with ZK rollups: Design, deployment, and evaluation of a real-time proof-of-concept, 2025, arXiv preprint arXiv:2506.00500.

[101] Cube, What is time to finality?, 2019, https://www.cube.exchange/what-is/time-to-finality?utm (Accessed 17 January 2026).

[102] D. Christopher, Privacy L2 Aztec is (almost) ready for primetime, 2025, https://www.bankless.com/read/privacy-l2-aztec-is-almost-ready-for-primetime (Accessed 15 January 2026).

[103] A. Foruj, Privacy L2 Aztec is (almost) ready for primetime, 2025, https://forum.aztec.network/t/how-are-apps-and-infras-affected-by-aztec-block-times-and-how-to-handle-re-orgs/7693 (Accessed 17 January 2026).

[104] Aztec, Security of the Aztec network: Audits of bigfield, 2025, https://aztec.network/blog/security-of-the-aztec-network-audits-of-bigfield?utm (Accessed 17 January 2026).

[105] Artsyl, Manual invoice processing vs. Automated invoice processing: A comprehensive comparison, 2025, https://www.artsyltech.com/manual-invoice-processing-vs-automated-invoice-processing (Accessed 17 January 2026).

[106] Resolve, 13 statistics that quantify cost per invoice in manual vs automated flows, 2025, https://resolvepay.com/blog/13-statistics-that-quantify-cost-per-invoice-in-manual-vs-automated-flows (Accessed 18 January 2026).

[107] W.C. Rivenbark, K.M. FitzGerald, S.H. Schelin, W.H. Yates, T. Runkle, Information Technology Investments: Metrics for Business Decisions, Institute of Government, University of North Carolina at Chapel Hill, 2001.

[108] Aztec, Dollars and sense: Cheap privacy with Atzec connect, 2025, https://aztec.network/blog/dollars-and-sense-cheap-privacy-with-aztec-connect?utm (Accessed 20 January 2026).

[109] YCHARTS, Ethereum average gas price (i:EGPND), 2025, https://ycharts.com/indicators/ethereum_average_gas_price?utm (Accessed 20 January 2026).

[110] M. El-Hajj, B. Oude Roelink, Evaluating the efficiency of zk-snark, zk-stark, and bulletproof in real-world scenarios: A benchmark study, Information 15 (8) (2024) 1–43, https://doi.org/10.3390/info15080463.