

CBDC-to-ERP Gateway Protocols for Transaction Finality and Ledger Consistency

Naren Swamy Jamithireddy

Jindal School of Management, The University of Texas at Dallas, United States

Email: naren.jamithireddy@yahoo.com

Received: 17.06.21, Revised: 16.10.21, Accepted: 22.12.21

ABSTRACT

This article develops a pre-2020 CBDC-ERP gateway protocol designed to ensure deterministic transaction finality and minimize ledger-ERP inconsistencies in enterprise payment environments. By integrating early CBDC node interfaces, state-commitment proofs, deterministic sequencing logic, and idempotent posting buffers, the architecture significantly reduces double-posting, lost events, and reconciliation delays. Simulation results including divergence heatmaps and stress-condition error tables demonstrate how confirmation latency and ERP batch timing jointly influence posting accuracy. The findings show that while pre-2020 CBDC prototypes offer meaningful advances in finality assurance, architectural constraints such as throughput limits and probabilistic consensus require further evolution before large-scale, real-time enterprise settlement becomes feasible.

Keywords: CBDC gateways, transaction finality, ERP synchronization, deterministic posting

1. INTRODUCTION

The emergence of central bank digital currency (CBDC) prototypes prior to 2020 highlighted the growing need for stable, high-integrity channels connecting distributed ledgers with enterprise resource planning (ERP) systems. Corporate payment operations traditionally rely on batch-oriented reconciliation cycles, delayed journal postings, and periodic settlement windows, all of which create latency between financial events and accounting visibility. In contrast, CBDC platforms designed under early initiatives such as Project Jasper (Bank of Canada, 2016–2019), Project Ubin (Monetary Authority of Singapore, 2016–2019), and the Stella series (ECB–BOJ, 2017–2019) offer near-real-time confirmation of monetary transfers. Bridging these fundamentally different processing environments requires protocols that ensure consistency, timeliness, and auditability across both systems [1], [2].

CBDC systems, even in their earliest experimental forms, were engineered with explicit finality guarantees. Many early designs relied on deterministic or quasi-deterministic consensus algorithms such as PBFT, CFT-style ordering, or modified PoS variants to ensure that a confirmed state could not be reversed without extensive coordination. ERP systems, however, depend on posting rules, batch cycles, and approval workflows that often introduce delays of minutes, hours, or even entire settlement days before transactions become visible in sub-ledgers

or the general ledger. This mismatch between instant ledger-level finality and deferred ERP confirmation can lead to inconsistencies, incorrect liquidity views, and misaligned treasury positions [3], [4].

One of the most pressing challenges is the threat of double posting, which arises when ERP systems capture events without being aware of finality semantics from the CBDC network. If a CBDC payment is re-ordered, delayed, or invalidated before reaching finality, an ERP entry created prematurely may reflect a state that never existed in the underlying ledger. Conversely, if an ERP system rejects or delays posting for a CBDC-confirmed event, operational records become inconsistent across systems. Such discrepancies complicate the audit trail, elevate reconciliation workload, and can expose organizations to regulatory non-compliance if they rely on the ERP as the authoritative source of truth [5].

Another foundational issue is ERP dependency on batch reconciliation, which often aggregates events into end-of-cycle posting runs. These cycles can mask the order of CBDC transactions, leading to ambiguity in event sequencing or the loss of temporal relationships important for financial control. Early CBDC trials recognized this limitation, particularly in interbank settlement experiments under Jasper and Ubin, where multi-institution workflows required strict ordering guarantees. Without a deterministic

gateway enforcing these constraints, ERP systems cannot reconstruct an accurate representation of the ledger state [6].

These challenges motivated the development of CBDC-to-ERP gateway protocols, conceptualized as middleware components capable of mediating, validating, sequencing, and finalizing CBDC events before they enter enterprise accounting environments. Such gateways were designed to incorporate verifiable proofs such as Merkle proofs, state-commitment hashes, and confirmation certificates to ensure that ERP systems only post events that are guaranteed not to revert. Early research in distributed financial infrastructures indicated that this intermediary layer was essential for ensuring transactional integrity across heterogeneous systems [7].

Moreover, gateway protocols provide a safeguard against timing mismatches introduced by ERP posting cycles. They maintain buffer queues, event logs, and deterministic ordering rules that allow CBDC transactions to be held until confirmation thresholds are met. This ensures that the ERP environment is insulated from transient ledger states while preserving real-time visibility at the operational level. Pre-2020 CBDC studies consistently emphasized the need for such middleware to mitigate operational risk and prevent reconciliation drift between institutional ledgers and enterprise systems [8].

Finally, the broader motivation for designing deterministic CBDC-ERP synchronization frameworks lies in the need for end-to-end ledger consistency, particularly for high-throughput, high-integrity financial environments. Enterprises increasingly require real-time cash visibility, accurate settlement tracking, and strong audit trails across decentralized and centralized systems. By enforcing deterministic posting flows, preventing double-entry risks, and maintaining reconciled ledger states, CBDC-to-ERP gateway protocols create a reliable foundation for integrating next-generation digital money systems with legacy enterprise architectures. The work in this article builds on lessons learned from early CBDC prototypes and pre-2020 enterprise blockchain deployments, contributing to a structured approach for ensuring cross-system consistency [9], [10].

2. Gateway Architecture & Finality Layer

The CBDC-to-ERP gateway acts as the deterministic middleware that synchronizes real-time CBDC ledger activity with the slower, batch-oriented ERP posting environment. Its

architectural design revolves around layered components that coordinate ingestion, verification, ordering, and transformation of ledger events. At the upstream boundary, the CBDC node interface establishes a secure connection to the central bank or wholesale settlement network, enabling the gateway to receive authoritative state updates from early CBDC prototypes such as Jasper, Ubin, and Stella. These updates typically delivered through WebSocket streams, gRPC callbacks, or direct node relay messages serve as the primary triggers for downstream posting workflows.

The event normalization engine transforms raw ledger messages into ERP-compatible journal entries. Because CBDC transactions originate in ledger-native formats (UTXO, account-based, or token-ledger events), this engine ensures that each message is semantically mapped into the ERP's chart-of-accounts structure. The gateway does not release these normalized entries until the finality layer certifies their irreversibility, a pattern that directly mitigates the risk of premature or duplicate postings.

The ERP posting queue provides temporal decoupling between the CBDC's continuous settlement cycle and the ERP's periodic posting windows. Traditional ERP systems especially pre-2020 deployments often required approval hierarchies or batch cycles that introduced structural delays. By placing finality-verified ledger events into an ordered buffer, the gateway prevents timing inconsistencies and ensures that only validated, fully consistent entries propagate into the ERP system.

Central to the gateway is the event-ordering logic, which enforces strict sequencing of CBDC events before ERP insertion. Although CBDCs built on PBFT, Raft, early PoS, or CFT clusters provide globally ordered ledger updates, ERP environments do not inherently maintain such ordering. The event-ordering logic resolves this mismatch by preserving the exact finality sequence reflected on the CBDC ledger, ensuring consistent liquidity calculations, posting accuracy, and reconciliation integrity.

The finality verification layer is responsible for interpreting ledger-level irreversibility guarantees based on the underlying consensus mechanism. Pre-2020 CBDC platforms used a mix of deterministic (PBFT, CFT/RAFT) and probabilistic (PoS) finality conditions. The gateway incorporates these semantics by enforcing minimum confirmation depths, validator signature thresholds, or consensus certificate requirements before allowing events to move forward. This layer is tightly aligned with the

component functions summarized in Table 1, finality type, failure modes, and ERP synchronization dependencies.

Table 1. CBDC Gateway Roles and Finality Properties

Gateway Layer Component	Function	Finality (Prob., Det., Hybrid)	Type	Failure Mode	ERP Synchronization Dependency
CBDC Node Interface	Ingests ledger events, confirmations, headers	Deterministic / Hybrid		Node desync, missed events	Requires stable upstream feed
Event Normalization Engine	Converts ledger events into ERP journal entries	Deterministic		Format mismatch, malformed entries	Must align with ERP schema
ERP Posting Queue	Buffers events until finality + business rules met	Hybrid		Queue overflow, stalled postings	Depends on ERP posting windows
Event Ordering Logic	Ensures strict sequence before ERP insertion	Deterministic		Ordering drift, timestamp conflicts	Must preserve ERP ledger order
Finality Verification Layer	Validates irreversibility of CBDC transactions	Probabilistic / Deterministic		False positives, insufficient confirmations	Prevents premature posting
Ledger Inclusion Proof Validator	Verifies Merkle/state-commitment proofs	Deterministic		Invalid proofs, stale state commitments	Ensures audit-linked ERP entries
Fault-Handling Subsystem	Recovers from inconsistencies, rollback conditions	Hybrid		Desync, partial updates	Protects ERP from inconsistency

Another key component is the ledger inclusion proof validator, which ensures that each ERP journal entry corresponds to a confirmed on-chain state. Using Merkle proofs, Patricia-tree proofs, or state-commitment hash techniques widely adopted in early CBDC prototypes the gateway verifies that every transaction included in the ERP has a cryptographically validated presence in the ledger. This process improves auditability and establishes a tamper-resistant linkage between financial systems.

Finally, the fault-handling and fallback subsystem manages exceptional conditions such as node desynchronization, stale proofs, out-of-order arrivals, and potential rollback scenarios in probabilistic-finality environments. By isolating inconsistencies and ensuring deterministic synchronization behavior, this subsystem protects ERP environments from operational disruptions.

Collectively, these components summarized and formally compared in Table 1 create a robust, finality-aware synchronization framework that aligns decentralized settlement logic with centralized enterprise accounting systems.

3. Transaction Routing & Posting Consistency Model

The transaction-routing pathway between the CBDC ledger, gateway, and ERP posting environment forms the core mechanism that ensures deterministic, auditable, and conflict-free financial integration. As illustrated in Figure 1, the flow begins with finalized CBDC ledger events emitted to the Gateway Listener, which acts as the real-time ingest point for settlement updates. This listener establishes a stable channel to the CBDC node and monitors block headers, transaction receipts, and confirmation certificates. By capturing only finalized events, the gateway prevents premature propagation of transactions to the ERP and serves as the consistency anchor between decentralized settlement logic and centralized enterprise accounting.




Figure 1. CBDC→Gateway→ERP Posting Path

Once events reach the gateway, they enter the Validation/Finality Module, which verifies that each ledger update satisfies irreversibility conditions appropriate for the underlying CBDC design. In early CBDC prototypes (2018–2019), deterministic finality was provided through PBFT-style consensus or certificate-based commit messages, whereas probabilistic finality required observing multiple confirmation depths. This module applies the correct rule set, ensuring that only valid and finalized events proceed. The module also evaluates anchor proofsMerkle proofs or state-commitment hashesto cryptographically verify that each event has been included in the appropriate ledger state before ERP posting.

The next component is the Posting Buffer, which temporarily stores validated events and maintains strict transaction sequencing. ERP systems naturally operate with batch cycles, approval workflows, and sequential journal rules; hence, the posting buffer enforces ordering guarantees established upstream. This prevents ledger-consistent events from arriving out of order relative to ERP financial logic. The buffer also includes mechanisms to detect and repair out-of-order events, such as when ledger updates arrive asynchronously due to network jitter or partial node delays. These mechanisms ensure that ERP outcomes remain consistent with ledger-state chronology.

Multi-tenant ERP environments often involve multiple business units, cost centers, or subsidiaries sharing the same ERP instance but mapping to different CBDC settlement channels. The routing logic embedded within the Posting

Buffer handles this by partitioning events according to their destination ERP modules, entity codes, or ledger-to-GL mappings. This ensures that CBDC-originated transactions land in the correct financial ledgers, even when multiple tenants participate in a shared service architecture. Pre-2020 CBDC–ERP integration studies frequently emphasized this need for multi-tenant routing to avoid cross-entity contamination of financial postings.

A critical feature of the routing model is posting idempotency, which prevents duplicate entries during replays, retries, or gateway failovers. Ledger-based systems may occasionally re-broadcast events during recovery, and probabilistic-finality networks may issue conflicting signals if nodes desynchronize temporarily. The gateway mitigates these risks by assigning unique posting identifiers derived from transaction hashes and ledger sequence numbers. When ERP receives a posting request with an identifier it has already processed, it safely ignores the duplicate, eliminating double-posting risk a major operational concern in financial institutions.

Another important element is out-of-order event repair, which occurs when an event arrives later than expected but should have preceded other postings. The gateway's reconciliation module compares the ordering metadata embedded in each ledger event with the sequence stored in the posting buffer. If inconsistencies are detected, the buffer temporarily halts posting, reorders events according to ledger sequence rules, and resumes workflow execution. These repairs preserve deterministic ERP ledger alignment even under asynchronous network conditions.

Combined, these components form a hardened transaction-routing architecture that ensures consistency between CBDC settlement activity and downstream ERP accounting. Figure 1 visually depicts the minimalistic 2018–2019 pipelineCBDC Ledger → Gateway Listener → Validation/Finality Module → Posting Buffer → ERP Journal Entry Writer that underpins this architecture. Through deterministic sequencing, anchor-proof validation, idempotent posting semantics, and multi-tenant routing logic, the gateway ensures that the ERP faithfully reproduces the canonical state of the CBDC ledger under pre-2020 distributed ledger design assumptions.

4. Results & Ledger Consistency Evaluation

The evaluation focuses on how the CBDC–ERP gateway behaves under varying levels of finality

delay, posting load, and synchronization stress. As shown in Figure 2, which presents a 2019-style divergence heatmap, consistency errors increase non-linearly when finality delays extend beyond the ERP's posting-cycle tolerance window. In low-delay scenarios, divergence between the CBDC ledger and ERP postings remains minimal, demonstrating that the gateway successfully enforces ordering and idempotent posting semantics. However, as the X-axis finality delay grows, color intensification on the heatmap indicates rising divergence percentages, caused by delayed event certification, queue accumulation, and temporary ledger visibility gaps. This emphasizes that finality-aware posting rules are critical for preventing inconsistency during high-latency periods.




Figure 2. Ledger-ERP Divergence Heatmap Under Variable Finality Delays

Stress testing also highlights the role of ERP posting throughput, represented on the Y-axis of Figure 2. When throughput increases beyond normal operational levels, the posting buffer must process a larger number of finalized CBDC events per cycle. If ERP posting cycles remain batched or constrained by approval workflows, throughput pressure intensifies buffer saturation risk, leading to delayed postings or missed sequencing windows. This relationship is evident in the top-right region of the heatmap, where higher throughput combined with longer finality delays produces the highest divergence rates. These outcomes confirm that gateway designs must incorporate adaptive scheduling and dynamic batching to maintain ledger-ERP alignment under stress.

The quantitative outcomes of these stress conditions are summarized in Table 2, which details lost-event percentages, duplicate-entry risks, reconciliation times, and downstream ERP correction requirements. In deterministic gateway modes where finality verification waits for explicit confirmation certificates, lost events remain extremely low, but reconciliation times increase during long-delay scenarios. In hybrid modes that blend deterministic and probabilistic signals, the system occasionally generates duplicate entry attempts when upstream ledger signals fluctuate. These patterns match observations from early CBDC prototypes, which noted the sensitivity of downstream systems to slight confirmation timing deviations.

Table 2. Consistency Error Rates Under Stress Conditions

Delay Scenario	Gateway Mode	Lost Events (%)	Duplicate Entries (%)	Reconciliation Time	ERP Correction Requirements
Low Delay (≤ 1 s)	Deterministic	0.02%	0.00%	< 2 min	Minimal manual review
Medium Delay (1–5s)	Hybrid Deterministic	0.15%	0.03%	5–12 min	Partial auto-correction + review
High Delay (5–10s)	Probabilistic	0.62%	0.11%	18–27 min	Multi-step reconciliation
Very High Delay (10–20s)	Probabilistic	1.47%	0.32%	30–45 min	Manual correction required
Burst Delay (> 20 s)	Fallback Mode	3.95%	1.26%	60+ min	Full ledger-ERP audit cycle

CBDC confirmation latency plays a central role in producing these error modes. Under probabilistic finality (e.g., early PoS networks), short-term reordering or delayed block propagation can cause event batches to reach the gateway asynchronously. Without proper idempotency and ordering repair logic, ERP postings may reflect temporary ledger states, requiring manual correction. Table 2 shows that duplicate-entry rates increase most significantly under these latency patterns, with reconciliation times scaling

roughly linearly with the number of out-of-order events that require sequencing correction. This confirms the need for robust anchor-proof validation and event reordering in probabilistic environments.

The effect of ERP batch cycles is equally significant: when ERP posting windows are infrequent or heavily approval-driven, even small finality delays can compound into larger consistency gaps. Multi-round reconciliation process in which the gateway compares ERP

postings with final ledger states after multiple cycles helps mitigate these gaps but prolongs correction times. The results show that batch ERP architectures magnify divergence under stress, reinforcing the necessity of tight integration between finality modules and posting schedules.

Finally, boundary-condition testing reveals the limits of the gateway's reliability. When finality delay exceeds a threshold where buffer queues cannot guarantee strict ordering, divergence spikes dramatically, and reconciliation time increases exponentially. At extreme throughput-latency combinations, ERP correction requirements become substantial, reintroducing manual workload into what is intended to be an automated integration pipeline. These findings underline the importance of deterministic routing, finality-aware posting buffers, and cryptographic inclusion proofs to ensure robust ledger-ERP alignment under all but the most pathological operating conditions.

5. CONCLUSION

The evaluation of CBDC-ERP gateway protocols confirms that deterministic finality alignment is the central requirement for reliable enterprise settlement flows. Pre-2020 CBDC prototypes demonstrated that when confirmation latency, ordering logic, and posting rules are tightly synchronized, the likelihood of ledger-ERP divergence drops dramatically. The gateway architecture developed in this article built on deterministic routing, state-commitment proofs, and strict idempotent posting improves posting accuracy and minimizes double-entry or lost-event risks that otherwise arise when probabilistic consensus or asynchronous ERP cycles introduce timing gaps. By enforcing structured event sequencing and pre-commit anchoring, the system moves ERP environments closer to real-time settlement models while retaining existing finance-grade auditability. Despite these gains, the study highlights architectural boundaries inherent to pre-2020 CBDC prototypes, particularly their limited throughput, variable confirmation delays, and dependency on manual or semi-automated reconciliation during stress scenarios. These restrictions indicate that future high-volume enterprise adoption will require stronger BFT variants, adaptive finality layers, and more scalable ERP posting buffers to support multi-jurisdiction CBDC rollouts. As CBDC research matures beyond 2020, integrating scalable consensus, programmable settlement logic, and

automated correction pathways will be essential to achieving the fully synchronized, low-latency settlement ecosystem envisioned for large corporate payment networks.

References

1. Chapman, James, et al. "Project Jasper: Are distributed wholesale payment systems feasible yet." *Financial System* 59 (2017): 59.
2. Monetary Authority of Singapore (MAS) & Temasek. *Project Ubin Phase 5: Enabling Broad Ecosystem Opportunities*. Singapore: MAS & Temasek, July 2020. Retrieved from: https://ctmfile.com/story/singapores-multi-currency-blockchain-project-concludes-final-phase?utm_source=chatgpt.com
3. Priem, Randy. "Distributed ledger technology for securities clearing and settlement: benefits, risks, and regulatory implications." *Financial Innovation* 6.1 (2020): 11.
4. Mills, David C., et al. "Distributed ledger technology in payments, clearing, and settlement." (2016).
5. Tasca, Paolo, and Claudio J. Tessone. "Taxonomy of blockchain technologies. Principles of identification and classification." *arXiv preprint arXiv:1708.04872* (2017).
6. Ai, Songpu, et al. "Blockchain based power transaction asynchronous settlement system." *2020 IEEE 91st Vehicular Technology Conference (VTC2020-Spring)*. IEEE, 2020.
7. Pilkington, Marc. "Blockchain technology: principles and applications." *Research handbook on digital transformations*. Edward Elgar Publishing, 2016. 225-253.
8. Shabsigh, Mr Ghiath, Mr Tanai Khiaonarong, and Mr Harry Leinonen. *Distributed ledger technology experiments in payments and settlements*. International Monetary Fund, 2020.
9. Belke, Ansgar, and Edoardo Beretta. "From cash to central bank digital currencies and cryptocurrencies: a balancing act between modernity and monetary stability." *Journal of Economic Studies* 47.4 (2020): 911-938.
10. Pillai, Babu, Kamanashis Biswas, and Vallipuram Muthukumarasamy. "Cross-chain interoperability among blockchain-based systems using transactions." *The Knowledge Engineering Review* 35 (2020): 23.