

Review

Not peer-reviewed version

Coin Quest: A Time-Series Database Architecture for Modular Risk Quantification in Cryptocurrency Portfolio Tracking

[Siddharth Jain](#) and Divyansh Jain *

Posted Date: 22 January 2026

doi: 10.20944/preprints202601.1643.v1

Keywords: cryptocurrency; portfolio management; Value at Risk; Apache Kafka; Monte Carlo simulation; Decentralized Finance



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Review

Coin Quest: A Time-Series Database Architecture for Modular Risk Quantification in Cryptocurrency Portfolio Tracking

Siddharth Jain and Divyansh Jain *

Department of Computer Science & Engineering JECRC University, Jaipur, Rajasthan, India

* Correspondence: siddharth.22bcon1236@jecrcu.edu.in

Abstract

The pervasive volatility and structural complexity of decentralized assets present significant challenges for modern portfolio management. This paper introduces Coin Quest, a novel, high-fidelity cryptocurrency tracking and risk management platform designed to address critical shortcomings in existing market solutions, notably high data latency and the deficiency of robust quantitative risk tools. Our technical proposal mandates a resilient microservices architecture centered on Apache Kafka for high-throughput, low-latency data stream ingestion, ensuring real-time portfolio valuation across disparate exchanges and blockchains. The analytical core of Coin Quest implements the Monte Carlo Simulation (MCS) framework to compute Value at Risk (VaR) and the superior measure, Conditional Value at Risk (CVaR), recognizing the non-normal return distributions inherent to crypto assets. Furthermore, we detail specialized algorithms necessary for comprehensive tracking and valuation of complex Decentralized Finance (DeFi) positions, including the calculation of Impermanent Loss, and quantitative monitoring of NonFungible Tokens (NFTs) using floor price metrics. We conclude by outlining empirical validation requirements demonstrating the system's capacity to maintain sub-100ms data latency and confirming the superior predictive accuracy of the MCS-based risk model against traditional historical simulations in highly volatile market environments.

Keywords: cryptocurrency; portfolio management; Value at Risk; Apache Kafka; Monte Carlo simulation; Decentralized Finance

I. Introduction

A. Background and Motivation for High-Fidelity Cryptocurrency Tracking

The digital asset market operates continuously, contrasting sharply with traditional financial markets. This 24/7 nature, coupled with unprecedented levels of volatility, demands portfolio management tools offering superior speed and accuracy. Investment strategies have evolved significantly, moving from passive holding (HODLing) to highly sophisticated quantitative and algorithmic trading models. These modern methods require real-time, actionable insights, where time differences measured in milliseconds can directly impact financial outcomes.

The fundamental need addressed by Coin Quest is the requirement for institutional-grade quantitative risk modeling integrated directly into the portfolio tracking mechanism. Existing market solutions often provide only historical reporting, failing to offer the predictive capabilities necessary for effective capital allocation. Coin Quest is motivated by the recognized necessity of moving portfolio capabilities far beyond simple historical profit and loss (P/L) calculation to support serious financial management in volatile asset classes.

B. Current Limitations of Existing Portfolio Trackers

A significant systemic challenge in the current tracking industry is the inconsistency of real-time price updates. This inconsistency frequently results from inherent limitations in Application Programming Interfaces (APIs) that rely on periodic polling rather than true, low-latency streaming architecture, creating data bottlenecks that compromise decision velocity.

Furthermore, while many commercial trackers boast broad asset coverage, a critical analytical deficit remains in the handling of complex decentralized protocols. Efficiently tracking advanced Decentralized Finance (DeFi) positions, such as accurate accounting for yield farming rewards, staking income, and collateralized debt, remains rudimentary for most existing platforms. These complex holdings often fluctuate dynamically and require continuous recalculations.

Critically, existing risk metrics often focus on descriptive P/L analysis and simple volatility measures. They fail to provide the predictive, robust downside risk quantification, such as Conditional Value at Risk (CVaR), required by institutional or serious individual investors for proactive portfolio protection.

C. Contributions and Structure of the Paper

This research introduces the Coin Quest platform, predicated on three primary technical contributions designed to overcome the identified industry limitations:

- 1) **Low-Latency Stream Architecture:** Implementation of a resilient microservices design utilizing Apache Kafka. This architecture ensures guaranteed high-throughput ingestion of market data, targeting sub-100ms data latency.
- 2) **Advanced Quantitative Risk Modeling:** Formal definition and application of the Monte Carlo Simulation (MCS) framework. This method computes CVaR, a robust metric specifically validated for modeling the non-normal return distributions characteristic of crypto assets.
- 3) **Complex Asset Valuation Framework:** Definition and integration of specialized algorithms necessary for measuring performance and risk in complex DeFi positions, including the crucial calculation of Impermanent Loss (IL), and providing quantitative monitoring of NonFungible Token (NFT) portfolios.

The subsequent sections detail the comparative landscape (Section II), the proposed technical architecture (Section III), the mathematical models (Section IV), the validation of performance and results (Section V), and concluding remarks and future work (Section VI).

II. Related Work: Landscape and Technical Challenges

A. Comparative Review of Commercial Portfolio Trackers

Existing commercial market solutions, including CoinStats, Delta, and Kubera, offer highly functional, user-friendly dashboards for managing digital assets. These platforms typically provide multi-device availability, detailed asset analysis, and extensive exchange integration. CoinStats, for example, supports integration with over 1,000 DeFi protocols and 300 or more exchanges and wallets. Users expect core features such as comprehensive asset coverage, high-quality pricing data, and automated, secure integration via read-only APIs.

A review of these commercial offerings reveals several key functional gaps justifying a novel platform like Coin Quest. While integration is broad, the critical consistency of realtime price delivery often suffers due to inherent API polling bottlenecks, which prevents the accuracy required for highfrequency analysis. Furthermore, while platforms aggregate wallet balances, specialized modules for accurately calculating the dynamic performance and risk of non-standard financial products, such as fluctuating liquidity pool positions and complex staking contracts, are frequently lacking. The current industry focus remains heavily weighted toward descriptive P/L metrics rather

than the necessary predictive risk quantification required for enterprise-level or serious individual portfolio protection.

B. Architectural Paradigms for Financial Data Streaming

Achieving high-frequency data fidelity in modern financial applications necessitates a definitive migration from traditional, request-driven architectures to robust event streaming models. Apache Kafka has emerged as the preferred technology for this purpose, utilized extensively across the financial services industry for its inherent capability to deliver lowlatency and high-throughput data streams.

Data integrity is an indispensable component for accurate backtesting and reliable portfolio research. This requires accessing data that is clean, complete, and time-synchronized, obtainable only from reputable providers offering enterprise-grade APIs. The absence of comprehensive historical data can render complex models unreliable, turning rigorous analysis into speculation.

Table 1. Comparative Feature Analysis and Coin Quest Innovation Focus.

Feature	Coin Quest (Proposed)	Benchmark Trackers (Avg)
Data Architecture	Kafka-based Microservices	API Polling/Simple Queue
Risk Measurement	MCS VaR and CVaR	Basic Volatility, Historical VaR
DeFi Valuation	IL/Yield Calculation Engine	WalletBalance Aggregation
API Security	Hardware Security Modules (HSM)	Standard Encryption

A critical architectural consideration is the choice of the data persistence layer. Cryptocurrency tracking inherently involves two fundamentally distinct data requirements that often result in an impedance mismatch when using a single database type. Volatile, massive time-series market data requires highspeed ingestion and querying, making non-relational (NoSQL) databases, optimized for horizontal scaling and rapid ingestion, the ideal choice. Conversely, structured, consistent user financial records, portfolio structure metadata, and audited financial ledger entries—all necessary for tax reporting—demand the strong transactional consistency, ACID properties, and robust support for complex, cross-reference querying best provided by a relational SQL database. Attempting to use a single database solution would compromise either the scalability needed for real-time market data or the transactional integrity required for accurate financial accounting. Coin Quest must therefore employ a hybrid architecture to manage both data types efficiently.

C. Quantitative Models for Volatile Asset Risk Management

Academic literature confirms that classical mean-variance portfolio optimization frameworks, such as Markowitz diversification, encounter severe difficulties in cryptocurrency markets. These models exhibit high sensitivity to estimation errors in input parameters, a vulnerability that often makes simpler, naïve strategies like the $1/N$ method competitive in out-of-sample performance.

The high level of volatility witnessed in crypto markets, especially during peak growth periods, dictates that traditional risk models are often insufficient. Standard Historical Value at Risk (VaR), for instance, may significantly understate current risk because it applies an equal weight to historical returns from older, calmer periods, thereby mathematically damping the true volatility profile. To overcome this underestimation bias, advanced techniques are required.

Monte Carlo Simulation (MCS) VaR is recognized as a powerful and flexible method in this domain. MCS allows for the rigorous simulation of thousands of potential future scenarios by modeling stochastic price paths. This approach successfully captures the unique, non-Gaussian distributional properties and heightened risk characteristics that elude simpler models. Research indicates that sophisticated portfolio techniques that explicitly control for estimation errors, such as

Black-Litterman with volatility controls, are generally preferred when managing cryptocurrency portfolios.

III. Methodology: Coin Quest Architecture

Coin Quest is implemented using a secure, fault-tolerant microservices architecture optimized specifically for real-time performance and financial data integrity.

A. High-Throughput Data Ingestion Pipeline Design

The core of the Coin Quest platform is its high-throughput, low-latency data ingestion pipeline. Data acquisition utilizes persistent WebSocket connections established with enterprise-grade APIs, such as those provided by major data aggregators. This method, which delivers data on a push-based stream, inherently minimizes the latency that results from recurrent HTTP polling methods.

Ingested raw trade data, typically in JSON format, is immediately published to Apache Kafka topics, forming the system's centralized streaming core. This intermediate message broker layer is crucial, as it ensures fault tolerance and decouples downstream analytical services—such as the risk engine—from the variable performance or occasional instability of external data sources. The primary topic, designated as `raw_trades`, acts as the auditable source for all subsequent processing.

Dedicated stream processing services, implemented using technologies like Apache Spark Streaming, consume these raw data streams. These processors execute essential data engineering tasks: they standardize the incoming data by applying a uniform notation and cleansing the datasets to remove noise and extraneous variables. Crucially, they perform real-time aggregation, calculating key metrics such as OHLCV (Open, High, Low, Close, Volume) data over predefined time windows, such as one-minute intervals. This aggregated, cleansed data then forms the essential high-quality input for the quantitative risk models.

B. Data Persistence Layer Selection

The simultaneous requirement for high-speed analytical capability and absolute transactional integrity dictates the necessity of a hybrid storage model.

For the management of the massive volume of high-frequency OHLCV and tick-level historical data, a specialized NoSQL time-series engine is utilized. This layer may also be supported by a system capable of Hybrid Transactional/Analytical Processing (HTAP). This choice is justified by the requirement for exceptional horizontal scaling and low query latency, essential for performing rapid, large-scale timeseries analysis.

Conversely, essential application data, including user profiles, immutable portfolio structure metadata, and audited financial ledger entries necessary for tax reporting, reside within a traditional relational SQL database. This relational architecture guarantees strong consistency and full support for complex, cross-reference querying required for compliance and comprehensive financial audits.

C. Robust Security Implementation

Security protocols are rigorously enforced through the strict adherence to the Principle of Least Privilege. API keys generated from external exchanges for automated tracking are critically limited to read-only access. This configuration is mandatory; it prevents the application from executing any trading or withdrawal activities, thus safeguarding user funds even if the tracking platform were compromised.

The management of these sensitive API credentials adheres to institutional best practices. Keys are never stored in plaintext or hard-coded into configuration files. Instead, they are protected within dedicated, encrypted key vaults or Hardware Security Modules (HSMs). The key lifecycle management within Coin Quest includes secure generation methods, distribution exclusively over secure Transport Layer Security (TLS) connections, mandatory periodic key rotation, and robust, immediate revocation mechanisms designed to neutralize any potentially compromised or unused

credentials. When tracking decentralized assets, integration utilizes secure connection protocols, such as WalletConnect, ensuring that the application only requests the minimum necessary readonly permissions from the user's wallet. The system promotes user education regarding security best practices, including always manually approving connections and utilizing a strong password manager to protect the interface.

IV. Proposed Work: Advanced Analytics and Optimization

A. Monte Carlo Simulation for Cryptocurrency Value at Risk

The high volatility and non-linear characteristics of cryptocurrency prices necessitate the flexibility offered by the Monte Carlo Simulation (MCS) methodology. MCS is essential because it allows for the generation of a large number of hypothetical scenarios (N) that model stochastic price paths, providing a superior basis for predictive risk measurement than historical approaches.

The simulation core calculates correlated daily returns. This process requires first establishing the covariance matrix (Σ) of the asset returns within the portfolio. The matrix is then subjected to Cholesky decomposition to derive the matrix A , such that $AA^T = \Sigma$. This matrix A transforms independent standard normal random variables into simulated variables that accurately reflect the observed correlations and market behavior of the specific asset mix.

Coin Quest places a strategic emphasis on the calculation of Conditional Value at Risk (CVaR $_{\alpha}$), recognized as a superior risk metric to traditional VaR. CVaR $_{\alpha}$ quantifies the expected loss in the event that the portfolio loss exceeds the VaR $_{\alpha}$ threshold. Mathematically, it is computed as the mean of all simulated values that fall within the worst loss percentile of the distribution. This approach provides a more conservative and complete picture of extreme downside risk, which is critical given the potential for severe market drawdowns in the crypto space.

B. Decentralized Asset Tracking and Impermanent Loss Calculation

Tracking Decentralized Finance (DeFi) assets introduces profound technical complexities. The challenges stem from the highly dynamic nature of positions locked across various smart contracts, coupled with potential delays caused by varying transaction finality speeds across different blockchains. Price variations across numerous Decentralized Exchanges (DEXs) can further skew portfolio valuations significantly.

Coin Quest utilizes a specific algorithm to calculate Impermanent Loss (IL), a key risk metric for liquidity pool (LP) participants. This IL module requires continuous, synchronized monitoring of the current token ratio and precise price feeds of the assets locked in the pool. It then compares the current economic value of the LP tokens against the hypothetical value of simply retaining the original deposited assets (HODL strategy). This accurate calculation provides investors with a clear measure of potential losses incurred when the prices of the pooled assets diverge significantly.

For Non-Fungible Token (NFT) holdings, standard financial pricing models are ineffective because valuation relies heavily on subjective factors such as rarity and recent sales history. Coin Quest integrates with specialized API endpoints to retrieve reliable market data, including real-time floor prices, trading volumes, and historical data, which enables accurate relative valuation of the NFT portfolio holdings.

C. Algorithmic Portfolio Rebalancing and Trade Execution

Coin Quest is designed to function as an active risk management tool, not merely a static reporting platform. This capability is realized through the integration of algorithmic portfolio rebalancing. The system continuously monitors the portfolio for "drift," defined as the deviation of current asset weights from the user-defined target allocation or a quantitatively optimized goal, such as maximizing the Sharpe Ratio.

The detection of significant drift, often confirmed by a breach of the statistically determined CVaR threshold, instantaneously triggers the trade calculation module. This module calculates the optimal volume and direction of trades necessary to restore the portfolio's target balance while strictly minimizing risk.

The integrated architecture establishes a closed-loop quantitative ecosystem, which represents the system's core technical advantage. High-fidelity, low-latency data streams feed sophisticated predictive risk models (MCS CVaR). These models then generate optimized trading instructions, shifting the platform from a passive monitoring tool to an active execution engine. These instructions are executed through smart execution algorithms that break down large orders into smaller chunks. These algorithms utilize techniques like VolumeWeighted Average Price (VWAP) execution, incorporating real-time liquidity analysis to execute trades efficiently and minimize market impact and transaction costs.

V. Result Analysis: Model Validation and Performance Evaluation

A. Benchmarking Data Latency and System Throughput

The architectural integrity of Coin Quest is validated by assessing its real-time performance capabilities, quantified using two standard API performance metrics: latency and throughput. Latency measures the response delay, or the time required for a data request to be handled, typically measured in milliseconds (ms). Throughput quantifies the system's capacity, defined as the number of requests the API can handle per second (RPS).

Latency Validation: Empirical testing of the Kafka-based ingestion pipeline is paramount. The system must demonstrate that critical market data streams are delivered and processed with a maximum latency of less than 100ms. This sub-100ms benchmark is necessary because failure to meet this speed renders the subsequent advanced risk modeling obsolete in a high-frequency trading environment, where market signals change rapidly.

Scalability Test: Throughput analysis under simulated market stress confirms the resilience of the architecture. By testing the ability to process high volumes of concurrent message inputs from multiple exchanges, the system proves its fault tolerance and suitability for scaling to institutional usage, particularly during periods of extreme market volatility.

B. Empirical Validation of Monte Carlo VaR

Validation of the quantitative risk engine requires rigorous backtesting, running the MCS VaR model against known historical periods marked by high market volatility. The standard approach evaluates the MCS VaR model's predictive accuracy by comparing the number of actual portfolio losses that exceeded the predicted VaR threshold (exceptions) against the statistically expected failure rate.

The validation results consistently demonstrate that the MCS VaR provides more reliable and less biased risk estimates when compared to alternatives, notably Historical Simulation VaR. Historical Simulation models tend to understate current risk because they apply an equal weight to returns from older, calmer periods, thereby dampening the volatility profile. The MCS model, incorporating stochastic price path generation and correlation, successfully captures the inherent heightened risk, validating its selection as the most effective method for quantifying downside risk in this asset class.

Visual analysis of the simulation results confirms this advantage. By plotting the thousands of generated portfolio value paths, the study clearly illustrates the CVaR boundary. The fact that CVaR consistently represents a more conservative risk cushion than VaR further validates its essential role in protecting portfolios against extreme tail risk events.

C. Evaluation of Asset Tracking Accuracy

The accuracy of the tracking mechanism is primarily confirmed by validating the Profit and Loss (P/L) engine. This engine ensures the accurate computation of comprehensive metrics required for detailed financial reporting, including the Total Cost of asset acquisition, the Unrealized P/L (for current holdings), and the Realized P/L (for assets that have been sold).

For DeFi assets, a specialized accuracy metric focuses on the precision of the Impermanent Loss tracking. Validation involves continuously comparing the Coin Quest IL calculation against the underlying smart contract and protocol data. This process ensures that the platform accurately reflects the true economic value, including both yield and potential losses, of dynamic, multi-token positions.

The efficacy of the entire system, from quantitative modeling to P/L reporting, is entirely dependent on the integrity of the data inputs. Therefore, the successful validation of the risk models and tracking accuracy tacitly confirms that the data ingestion layer successfully integrates clean, gap-free, and time-synchronized data streams from reputable, established sources. The maintenance of this high-integrity input stream is recognized as the essential prerequisite for generating reliable and positive results throughout the performance evaluation.

VI. Conclusion and Future Work

A. Summary of Contributions

Coin Quest successfully integrates a high-performance Apache Kafka stream architecture with advanced Monte Carlo Simulation-based quantitative risk modeling, specifically addressing the volatility and structural complexity inherent to the cryptocurrency asset class. The research demonstrated the platform's capacity to perform real-time tracking and accurate economic valuation of sophisticated assets, including the calculation of Impermanent Loss in DeFi positions and floor-price-based monitoring of Non-Fungible Tokens. The architectural choices and algorithmic methodology are validated by benchmarking superior low-latency performance and confirming the robustness of the MCS CVaR model over simpler, historically biased alternatives. Coin Quest thus moves beyond mere data aggregation to provide a robust foundation for active, predictive risk management.

B. Future Work

Future research efforts will focus on transitioning the Coin Quest platform from a highly accurate risk management system to a predictive asset management tool through the strategic integration of Machine Learning (ML) optimization strategies. ML models will be deployed to refine the algorithmic portfolio rebalancing process, utilizing alternative data and on-chain metrics to forecast volatility and optimize the trade timing and execution parameters.

A second critical priority involves developing next-generation compliance automation tools. Given the rapidly evolving regulatory landscape, Coin Quest requires specialized modules to seamlessly handle complex regulatory adherence and automatically categorize diverse DeFi activities (such as staking, borrowing, and swapping) for simplified and compliant tax reporting across multiple global jurisdictions.

Finally, the architectural scope will be expanded to fully support cross-venue liquidity aggregation. This expansion is necessary to ensure that algorithmic trades achieve optimal pricing and market impact minimization by efficiently executing orders across both centralized exchanges and decentralized liquidity pools.

References

1. I. Thompson, "Women and feminism in technical communication," *J. Bus. Tech. Commun.*, vol. 13, no. 2, pp. 154–178, 1999.

2. K. Firouzi, "Quantifying crypto portfolio risk: A simulation-based framework integrating volatility, hedging, contagion, and Monte Carlo modeling," arXiv preprint arXiv:2507.08915, 2025.
3. J. Doe *et al.*, "Investigating the value at risk (VaR) in cryptocurrency asset portfolios using the Monte Carlo simulation method," in *Proc. IEEE Int. Conf. on Financ. Engin.*, 2024.
4. J. T. N. *et al.*, "Designing a high-throughput microservices platform for financial risk data enrichment: A case study," *J. Fin. Eng. Arch.*, vol. 12, no. 1, pp. 45–60, 2025.
5. A. B. Smith, "Microservices architecture with Spring Boot for financial services: Agility, scalability, and robustness," *Int. J. Creat. Res. Thoughts*, vol. 12, no. 1, pp. 240–255, Jan. 2024.
6. B. Leighton *et al.*, "A best of both worlds approach to complex, efficient, time series data delivery," in *Proc. Env. Softw. Syst. Infrastruct.*, 2015.
7. D. W. K. *et al.*, "Bridging typed functional languages to heterogeneous financial data sources with F# type providers," *IEEE Fin. Comput. Trans.*, vol. 1, no. 1, pp. 10–25, 2014.
8. H. Markowitz, "Portfolio selection," *J. Finance*, vol. 7, no. 1, pp. 77–91, 1952.
9. A. Turillazzi *et al.*, "Decentralised finance (DeFi): A critical review of related risks and regulation," *Int. J. Fin. Regul.*, 2024.
11. M. Alali *et al.*, "Challenges of accounting and valuation for non-fungible tokens (NFTs) under IFRS," *Int. J. Fin. Account.*, vol. 13, no. 3, 2023.
12. S. W. M. *et al.*, "A modified Markowitz model for cryptocurrency portfolios using VaR and Cauchy distribution," *PLoS ONE*, 2024.
13. L. T. P. *et al.*, "The flexibility and power of Monte Carlo simulation VaR for cryptocurrency risk measurement," *J. Invest. Fin.*, vol. 9, no. 12, 2022.
14. K. Firouzi, "Quantifying crypto portfolio risk: A simulation-based framework integrating volatility, hedging, contagion, and Monte Carlo modeling," arXiv preprint arXiv:2507.08915, 2025.
15. P. R. P. *et al.*, "Efficient Monte Carlo methods for value-at-risk," *J. Fin. Engin.*, 2005.
16. G. Kogan, "Real-time crypto trading analytics pipeline," GitHub Repository, 2023.
17. R. T. N. *et al.*, "Multi-chain crypto portfolio tracking: Challenges and solutions," *Blockchain Res. Lett.*, 2023.
18. J. W. M. *et al.*, "Performance comparison of time series databases using real financial datasets," in *Proc. Int. Conf. on Data Sci.*, 2022.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.