

# Principles of Building Event-Driven Architectures for Logistics Process Management

Dhaval Shah

Senior Software Engineer, Redwood City, California, USA

**Abstract:** *The article explores how to design event-driven architectures (EDAs) for logistics process management that can integrate heterogeneous IoT/IT signals into a governed event fabric, detect operational meaning via complex event processing (CEP), and preserve trust through selective, integrity-protected logging of “high-value” events rather than raw sensor chatter. This article draws on both hands-on implementation experience and performance-related lessons from CEP research, especially around how to handle complex queries using time and space indexing. The core outcome is a set of practical design principles, centered on a two-layer architecture. Drawing on implementation experience and recent CEP performance research, the article formulates a set of design principles for scalable and trustworthy logistics EDAs. The proposed two-layer architecture separates low-latency event handling from integrity-focused logging, enabling both real-time responsiveness and auditability. The contribution lies in translating CEP, blockchain, and event-time semantics into actionable architectural guidance for logistics process management.*

**Keywords:** event-driven architecture (EDA); complex event processing (CEP); low-code integration (Node-RED); world state database (CouchDB); Hierarchical Temporal Indexing (HTI).

## 1. Introduction

Modern logistics operations are highly complex and time-sensitive. Ongoing disruptions in supply chains- like backups at ports or delivery delays- have made it clear how important it is for logistics systems to be both fast and adaptable. One way to support them is through event-driven systems. These setups work by tracking what is happening in real time- like new orders, sensor alerts, or shipment changes- and reacting right away. This kind of design helps teams monitor their operations continuously and respond to changes, which is key in logistics, where timing directly affects service. Tools like Complex Event Processing (CEP) and sensor networks make this possible by spotting patterns in data streams and flagging issues early.

There has been growing interest in using event-based systems in logistics. In one case, a system for port operations used real-time data to adjust schedules. In supply chains, real-time analytics have also been used to track disruptions and respond earlier [10]. This paper examines how event-driven systems are applied in logistics and identifies design choices that shape effective real-time processing, scalability, and integration with AI and IoT technologies.

## 2. Methods and Materials

This work is based on a review of publications that examine aspects of event-driven architecture, artificial intelligence, and process management in logistics and supply chain contexts. Abbasi et al. proposed a novel indexing strategy to improve CEP database performance, yielding faster event query times [1]. Abril et al. developed an integrated event-driven, real-time tactical-operational optimization framework for smart port operations and showed substantial improvements in port planning [2]. Daios et al. conducted a comprehensive survey of artificial intelligence applications in supply chain management, covering key processes and identifying AI’s benefits and adoption challenges [3]. Du Plessis et al. [4] gathered a wide range of examples where AI

is being tried in freight logistics. They pointed out that while many of these use cases are still early, there is clear potential for AI to support better sustainability choices and smarter decisions overall. Ferreira et al. [5] went through recent research on automation and AI in logistics. Their review pointed out that robotics and machine learning tools can help improve logistics performance. Yet they also noted that in many cases, current practices still lag behind what the technology can actually do.

Hangl et al. [6] looked at a large set of earlier literature reviews- 44 in total- focused on AI adoption in supply chains. Matenga & Mpofu presented a blockchain-based cloud manufacturing case study for transport equipment production, implementing an event-driven microservices architecture (with blockchain integration) in a railcar manufacturing supply chain, achieving real-time part traceability and improved inventory control [7]. Oncioiu et al. [8] looked into how AI might be used in circular logistics, especially in contexts where tech, company structure, and environmental issues all come into play. Ramos Gutiérrez et al. [9], meanwhile, did a broad review—169 papers total—on business process management and event-based systems in logistics. Riad et al. [10] explored how AI can help make supply chains more resilient. Rosa-Bilbao et al. [11] built a system called CEPEDALoCo. It combines IoT-based event detection with blockchain logging, using a low-code platform to track complex events in real time and store them securely. Toth et al. [12] proposed a theoretical setup for using AI in logistics, combining knowledge management and adaptable processes. Their idea is that blending those two areas helps organizations succeed with AI tools in real-world logistics work.

Even with all the recent work, there are still important gaps in the literature. One issue is that many AI and optimization efforts focus on narrow problems and do not connect well across the supply chain. For example, Abril et al. [2] point out that most research on port operations stays focused on one area at a time—like terminals or yard activities—without linking them together. Another problem is that real-time data

is still underused. A lot of existing models rely on static data and can not react quickly to disruptions, mostly because they are not wired into live sensor or operations streams [2][9]. Some CEP engines have cannot process fast, unpredictable data at scale [1]. Many organizations hesitate to commit, either because they do not have the right skills in-house or because it is challenging to make the case for the investment [6]. This paper takes those problems as its starting point. It draws from recent studies to offer grounded ideas on how event-driven systems for logistics can be built and deployed.

### 3. Results and Discussion

Building effective event-driven architectures for logistics process management requires adhering to several core principles derived from recent research findings. These principles are discussed below, along with their practical implications and supporting examples from the literature.

One of the basics in event-driven logistics is responding in real time. In a port setting, for example, berth or crane schedules can be changed based on delays instead of sticking to a fixed plan. Abril et al. [2] tested this idea and showed that using events to drive decisions helped reallocate resources and keep operations running more smoothly. Control towers with AI components work similarly by watching live streams and making adjustments before problems escalate [10]. The architecture has to handle fast data and avoid lag. If a temperature spike in a cold-chain unit or a sudden wave of orders comes in, the system needs to process it right away and take action—whether by checking a rule set or running a predictive model. These systems move from tracking what happened to taking action before a malfunction. Figure 1 depicts a real-time event-driven control loop where heterogeneous logistics signals are turned into operational decisions fast enough to matter.

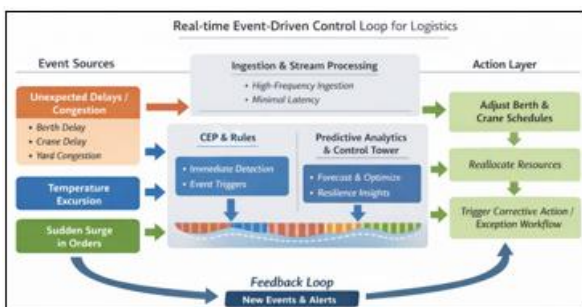


Figure 1. Real-time event-driven control loop for logistics EDAs

On the left of Figure 1, port disruption events (unexpected delays/congestion expressed as berth delay, crane delay, and yard congestion), cold-chain temperature excursions, and demand surges enter an ingestion & stream-processing layer designed for high-frequency ingestion and minimal latency; these streams are then evaluated in a dual decision lane- CEP & rules for immediate event detection and rule-based triggers, and predictive analytics / control-tower logic that forecasts impacts and supports resilience-oriented optimization- before flowing into the action layer (right), where the system adjusts berth and crane schedules, reallocates resources, and launches a corrective-action / exception workflow; crucially, the large feedback arrow emphasizes the EDA principle that actions

themselves emit new events & alerts, closing the loop and enabling continuous, proactive coordination rather than one-off, reactive firefighting.

Logistics systems deal with a flood of fast, varied data from both physical and digital sources. This includes GPS signals for location and speed, scan checkpoints (like RFID arrivals or departures), environmental readings (temperature, humidity, CO<sub>2</sub>), and system updates from platforms like WMS or TMS. An event-driven setup turns the inputs into a consistent stream so later steps in the process are clear. Rosa-Bilbao et al. [11] make this integration challenge concrete by showing how a low-code layer can ingest and harmonize measurements from 12 IoT devices, including environmental sensors (e.g., temperature, humidity, CO<sub>2</sub>) as well as three Netatmo weather stations, and then forward the cleaned stream into a CEP engine for real-time pattern detection [11]. Their implementation is not merely “multi-sensor ingestion” in the abstract: it operationalizes the integration principle as a pipeline of adapters and transformations (data acquisition → normalization → event emission) that converts vendor-specific payloads into a single event model suitable for continuous queries. Abril et al. [2] point to the same requirement from a different operational angle: adaptive port optimization depends on reliable connectivity to live feeds (IoT sensors, GPS tracking, weather), because the optimization layer can only re-plan when the underlying event stream reflects the current state of the terminal and its constraints [2]. Practically, this implies designing an ingestion tier that supports protocol mediation (e.g., IoT gateways and APIs), buffering, and schema governance so that new devices or message formats do not break downstream consumers. It also implies explicit treatment of time semantics (event-time vs. processing-time) and deduplication/idempotency, since logistics data routinely arrive late, out of order, or duplicated. When these concerns are handled systematically, heterogeneous inputs stop being isolated “signals” and instead become a shared operational substrate: events in one domain (e.g., a temperature excursion or a berth delay) can trigger coordinated responses in another (e.g., exception handling, re-routing, or service-level alerts), reducing information silos and enabling cross-stakeholder automation. Figure 2 formalizes the manuscript’s canonical event contract idea as the governance spine of an event-driven logistics stack.

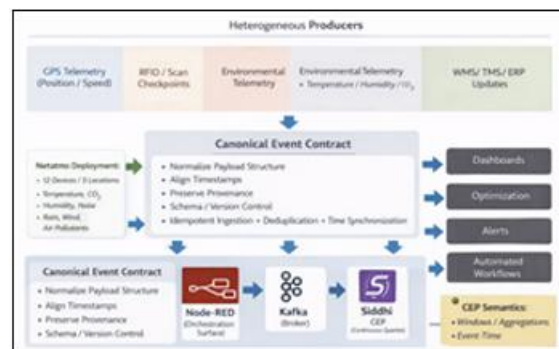


Figure 2. Canonical event contract pipeline for heterogeneous logistics signals

At the top of Figure 2, heterogeneous producers—GPS telemetry (position/speed), RFID/scan checkpoints, environmental telemetry (temperature/humidity/CO<sub>2</sub>), and

WMS/TMS/ERP updates- emit incompatible signals that must be made interoperable before any CEP can be trusted. The central contract layer shows the exact “uniqueness lexis” operations that make this possible: normalize payload structure, align timestamps, preserve provenance, enforce schema/version control, and ensure idempotent ingestion + deduplication + time synchronization so repeated or late arrivals don’t corrupt downstream reasoning. The left side of the figure shows a real-world setup using Netatmo sensors- 12 devices placed across three locations. They track temperature, humidity, CO<sub>2</sub>, noise, wind, rain, and some air pollutants. This setup shows why a shared data contract matters: the system needs to handle a variety of sensor types and local conditions. The middle of the figure shows the toolchain. Node-RED handles the orchestration, Kafka passes messages between components, and Siddhi runs continuous event queries based on timestamps and rolling windows. On the right, these outputs are used by dashboards, optimization tools, alerts, or automated workflows. This way the system turns incoming sensor data into something usable and structured.

Once event-driven systems start making decisions and triggering actions- rather than just monitoring- there is a stronger need for trust and accountability in how those decisions are tracked. One way to handle this is by saving a limited set of important events to a secure ledger, like a blockchain, where they can not be changed after the fact. Matenga and Mpofo [7] built this into a manufacturing setup using a permissioned blockchain to log production and transfer events. Rosa-Bilbao et al. [11] followed a similar idea. When their system detects a significant event, it logs the output—with time and context—directly to the blockchain. In that setup, the system logged more than 11,000 complex events. That shows this kind of high-volume recording can actually work—and still leave room for audit or review in case of a mistake. From a design point of view, logging should not be an afterthought. Often this means saving metadata or hashed versions instead of full sensor streams. In EDA terms, the ledger becomes a shared system-of-record for inter-organizational checkpoints: stakeholders can confirm that conditions were met (e.g., temperature remained within limits), reconstruct a sequence of responsibility handoffs, and validate why an automated decision was triggered. Even when blockchain is not used, the same principle applies via equivalently hardened mechanisms (append-only logs, WORM storage, cryptographic signing, and strict access controls): the event stream should remain actionable in real time, while a selected subset of events remains auditable over time. This combination strengthens the EDA’s role as both an operational “nervous system” and a credible source of truth for governance-heavy logistics processes.

Logistics systems produce a constant stream of data from both digital and physical sources. This includes vehicle tracking (GPS), identity scans (RFID), warehouse environment readings (temperature, humidity, CO<sub>2</sub>), and transactional updates from internal platforms like WMS or TMS. A working event-driven architecture needs more than just connectors between systems- it has to turn this input into a single, usable event stream. That means aligning timestamps, cleaning up payloads, and keeping track of where the data came from. Rosa-Bilbao et al. [11] built a low-code

integration layer that does this using Netatmo sensors installed across three sites. Their system pulls readings for temperature, humidity, noise, wind, and air quality, standardizes them, and feeds the results into a real-time pattern detection engine. Data moves step by step—from ingestion, through normalization, to event processing. Node-RED handles the orchestration, Kafka moves the messages between parts, and the CEP engine (Siddhi) runs continuous queries on the stream [11]. Different parts of the setup handle different functions. The system first takes incoming device data and puts it into a standard format so it can be used downstream. In practice, this integration principle implies interoperable protocols and middleware (IoT gateways, message brokers, APIs) plus governance mechanisms—schema/version control, idempotent ingestion, deduplication, and time synchronization—so that “one more device” or “one more data feed” does not destabilize the event fabric. When executed well, the EDA functions as an operational nervous system: it breaks silos by making cross-domain events visible under a shared model, enabling a condition detected in one area (e.g., an environmental excursion) to trigger coordinated responses elsewhere (e.g., exception workflows, re-routing, or service-level alerts) across stakeholders and systems.

As event-driven systems begin to take on more control-moving from monitoring into triggering actions and enforcing rules- the question of trust becomes harder to ignore. It is no longer enough to process events in real time; some of those decisions need to be logged in a way that can not be changed later. One approach is to combine automated control with blockchain-based logging, especially in settings where high-stakes events and shared accountability matter. Matenga and Mpofo [7] show how this works in a manufacturing context. They use a permissioned blockchain (Hyperledger Fabric) to record part production and transfer activity across different organizations. That setup allows for traceability and audit support, especially where fraud or contract violations are a concern. What is unusual in their design is that the blockchain is not storage: smart contracts and event-driven services can block incomplete transactions and mark suspicious ones. They also connect the ledger to CouchDB to store the world state, which can be analyzed separately to track performance or reliability [7]. Rosa-Bilbao et al. [11] take a similar route. Their system logs each detected complex event to the blockchain, with timestamps and context attached. Instead of saving raw sensor data, they focus on confirmed CEP outputs-over 11,000 of them in one deployment—which shows that large-scale logging like this is technically realistic. When multiple parties are involved, having an unchangeable record makes it easier to check what happened and when. This kind of log also helps confirm whether agreed conditions were met- for example, keeping goods below a certain temperature- and track handoffs between organizations.

From a design perspective, it is reasonable to figure out which types of events are most important for compliance or disputes. Those events, or hashed versions of them, can be written into a protected log that doesn’t allow edits. Meanwhile, the live event stream stays fast and responsive, with the blockchain layer acting as a slower but trusted record of key checkpoints. Done properly, secure event logging strengthens the credibility of the EDA’s decisions and establishes a verifiable operational history that stakeholders can independently audit-

turning the architecture into both a coordination engine and a trusted record of process execution. Figure 3 is meant to communicate a two-layer trust stack for logistics EDAs.



**Figure 3:** Secure event logging pattern in logistics EDAs: CEP outputs + permissioned ledger

On Figure 3 Tier A (Real-Time Event Fabric) is the responsiveness-optimized layer where CEP detects and acts on patterns, but—crucially—what gets persisted is not raw sensor chatter, rather detected complex events that are timestamped and contextualized, shown with the scale marker (11,000+ complex events) to emphasize that “trust” must work at volume, not just in toy cases. These confirmed events are then sent to a second layer—the integrity-protected ledger. This part of the system uses a permissioned blockchain (Hyperledger Fabric) to log key production and transfer records in a way that can not be altered later. In some setups, smart contracts and event-driven services can stop certain actions from going through if something looks wrong or trigger alerts for fraud. This layer also links to CouchDB, which holds the system state in a format that can be checked later for performance or reliability—without putting the ledger’s integrity at risk. There is also a small routing table that defines which kinds of events are logged on-chain. In this setup, only certain categories—like custody transfers, rule violations, contractual steps, or events likely to cause disputes—are stored as blockchain entries or hashed records. That helps keep auditability high without adding unnecessary cost.

Researchers have started trying AI methods in logistics setups that depend on event streams, but the results are still mixed. Some use machine learning or optimization tools to react faster or predict what might go wrong [3][10]. Riad et al. [10] give one case where the system looked ahead for possible problems and adjusted the plan—though how well that works likely depends on the context. Incorporating such AI components into an EDA means that incoming events are not only reacted to based on static rules, but can also be analyzed for patterns and fed into predictive models. Ferreira et al. [5] found that AI-enabled automation (including intelligent robotics and adaptive algorithms) can significantly boost operational efficiency and accuracy in logistics. In practical terms, this principle involves embedding machine learning modules or decision support engines within the event flow. For instance, streaming sensor and order data might be continuously fed into a trained model that predicts order fulfillment delays; if a high-risk prediction is made, the system generates a preventive event (e.g., alerting a manager or reassigning resources). Over time, AI algorithms can learn from the event history to improve predictions (creating a feedback loop that enhances performance). The goal is for the event-driven architecture to not just respond to what has

happened, but to anticipate what will happen, and to optimize logistics operations accordingly. This moves the system toward a self-optimizing state. When properly integrated, AI-driven analytics within an EDA can yield a more intelligent, autonomous logistics process that handles complexity and variability more effectively than rule-based automation alone [12].

An event-driven architecture should encompass the entire logistics process rather than optimize isolated segments. Instead of silos (where each department or stage operates separately), the EDA must integrate events across procurement, production, warehousing, and distribution so that all these activities are coordinated in real time. Abril et al. [2] argue that a holistic, enterprise-wide framework (spanning terminal operations to hinterland transport in a port, for example) is needed to avoid fragmented optimization. Ramos Gutiérrez et al. [9] observed that most EDA projects in logistics tend to focus on specific segments rather than the full process chain. As a result, it is often unclear how events in one part might affect other areas downstream. Without a structure for connecting these points, many systems miss the chance to adjust early. Designing for that kind of responsiveness across phases remains a challenge.

Finally, beyond technical design, the success of an event-driven architecture in logistics depends on organizational readiness and governance practices. Introducing a highly automated, data-driven system represents a significant change, and companies must ensure that their culture and processes adapt accordingly. Hangl et al. [6] stress that early change management—informing and involving the organization about the new technology—is critical to overcoming resistance and skill gaps. Training employees to work with real-time data dashboards and to trust (but verify) AI-generated decisions is an important aspect of this readiness. Governance frameworks should be established to define how data quality is maintained, how event-driven decisions are audited, and who is accountable in cases of system-triggered actions. Toth et al. [12] suggest that companies develop both strong data-driven knowledge practices and agile reconfiguration capabilities in parallel to fully realize the benefits of AI-driven, event-based systems. Some organizations have made progress by introducing dashboards or low-code rule editors [11], which allow staff with domain knowledge—not just developers—to adjust rules and interact with the system directly. Interfaces like these reduce the need for deep technical skills and make it easier to align automation with actual workflows. In summary, ensuring organizational readiness—through training, clear governance, and cultural buy-in—is a principle just as vital as any technical requirement. A well-governed, human-aware implementation will sustain the event-driven architecture’s performance and facilitate continuous improvement over time.

#### 4. Conclusion

Event-driven methods are proving useful for managing the pace and variability of modern logistics. Some systems that work with real-time event streams seem to support faster decisions, at least in specific use cases. There is some indication that automated responses can limit delays,

especially in unexpected situations. In one reported case, a port deployment showed shorter turnaround times and fewer delays with scheduling [2]. A similar architectural setup was applied in manufacturing, where tracking along production lines improved noticeably [7]. Forecasting gains were also observed in a study on multi-tier supply chains [10], though those findings likely depend on the specifics of data quality and system configuration. Because modular architectures tend to evolve incrementally, they can absorb new inputs—whether sensors, models, or rules—without forcing full-system redesigns.

Work on event-driven logistics systems still tends to concentrate on isolated components rather than on full process chains. One unresolved problem is the absence of frameworks that span procurement, operations, warehousing, and distribution in a single architectural model, as reported by Ramos Gutiérrez et al. [9]. Deployment practice is also uneven: Rosa-Bilbao et al. [11] show that CEP can be embedded into low-code environments, yet most implementations remain technically demanding and poorly aligned with operational roles. Hybrid event-processing setups, especially those combining different CEP styles, are still mostly explored in controlled or lab-based environments [1]. Even the models meant to improve resilience or optimization have not been put through sufficient testing in real logistics systems. Technical design alone is rarely enough. Many organizations face uncertainty about the tools themselves, and in some cases, lack staff who can implement or maintain them [6]. Environmental aims such as circular logistics and sustainability integration are widely discussed, yet they are rarely embedded in design work from the outset [4][8]. By articulating architectural principles that balance responsiveness, governance, and scalability, this study contributes a practical yet theoretically grounded foundation for future research and deployment of event-driven logistics systems.

## References

- [1] Abbasi M, Saad C, Elmeslahi O, Gui J. Optimizing database performance in complex event processing systems through indexing strategies. *Data* 2024;9(8):1-24. doi:10.3390/data9080093
- [2] Abril D, Paternina-Arboleda CD, Velasquez-Bermudez J. An integrated event-driven real-time tactical–operational optimization framework for smart port operations planning. *Logistics*. 2024;8(3):1-37. doi:10.3390/logistics8030065
- [3] Daios A, Siose A, Anastasopoulos D, Maselis G, Papadopoulos T. AI applications in supply chain management: A survey. *Appl Sci*. 2025;15(5):1-13. doi:10.3390/app15052775
- [4] Du Plessis MJ, Botes A, Boonzaaier C. Shaping the future of freight logistics: Use cases for integrating artificial intelligence and green transportation. *Sustainability*. 2025;17(4):1-18. doi:10.3390/su17041355
- [5] Ferreira B, Ramos E, Fontes D. A systematic literature review on the application of artificial intelligence and automation in logistics. *Logistics*. 2023;7(4):1-17. doi:10.3390/logistics7040080
- [6] Hangl J, Behrens VJ, Krause S. Barriers, drivers, and social considerations for AI adoption in supply chain management: A tertiary study. *Logistics*. 2022;6(3):1-22. doi:10.3390/logistics6030063
- [7] Matenga AE, Mpofo K. Blockchain-based cloud manufacturing SCM system for collaborative enterprise manufacturing: A case study of transport manufacturing. *Appl Sci*. 2022;12(17):1-18. doi:10.3390/app12178664
- [8] Oncioiu I, Nițescu DC, Petrescu AG, Dănescu T. Artificial intelligence-enabled digital transformation in circular logistics. *Logistics*. 2025;9(3):1-28. doi:10.3390/logistics9030102
- [9] Ramos Gutiérrez B, Reina Quintero AM, Parody L, Gómez López MT. When business processes meet complex events in logistics: A systematic mapping study. *Comput Ind*. 2023;144:1-16. doi:10.1016/j.compind.2022.103788
- [10] Riad M, Naimi M, Okar C. Enhancing supply chain resilience through artificial intelligence: Developing a comprehensive conceptual framework for AI implementation and supply chain optimization. *Logistics*. 2024;8(4):1-26. doi:10.3390/logistics8040111
- [11] Rosa-Bilbao J, Boubeta-Puig J, Rutle A. CEPEDALoCo: An event-driven architecture for integrating complex event processing and blockchain through low-code. *Internet Things*. 2023;22:1-16. doi:10.1016/j.iot.2023.100802
- [12] Toth Z, Alzamel SM, Erkan TE. AI integration in fundamental logistics components: Advanced theoretical framework for knowledge process capabilities and dynamic capabilities hybridization. *Logistics*. 2025;9(4):1-22. doi:10.3390/logistics9040140