

## AI-Enabled Multi-Agent Approach for Reliable and High-Performance Event Stream Management

**Arjun Sharma**

Department of Artificial Intelligence,  
National Institute of Digital Technology India

**Priya Mehta**

Department of Computer Science and  
AI, Institute of Advanced Computing India

**VOLUME 03 ISSUE 08 (2026)**

Published Date: 20 August 2026 // Page no.: - 29-34

---

### ABSTRACT

The increasing volume, velocity, and heterogeneity of event streams require intelligent management mechanisms capable of maintaining throughput, responsiveness, scalability, and operational reliability under continuously changing workloads. Conventional stream-processing architectures generally rely on predetermined routing, scheduling, and resource-management policies, which can become inefficient when event arrival rates, computational requirements, and system conditions change dynamically. This paper proposes a conceptual AI-enabled multi-agent approach for reliable and high-performance event stream management, integrating symbolic planning, learned representations, policy generation, temporal reasoning, and decentralized agent coordination. The theoretical foundation is derived exclusively from the supplied literature on planning, generalized policies, first-order representations, sketches, temporal logic, and description-logic state features. The proposed architecture models stream-management functions as interacting intelligent agents responsible for event observation, workload characterization, resource allocation, routing, fault adaptation, and policy refinement. The analysis indicates that combining learned state representations with generalized planning policies can improve adaptability while reducing dependence on manually specified control rules. Temporal reasoning further provides a basis for deadline-aware and sequence-sensitive stream decisions, whereas planning sketches can constrain decision complexity in large operational spaces. The study positions multi-agent event stream management as a synthesis of symbolic planning and adaptive AI rather than as a purely reactive stream-processing mechanism. The principal contribution is a research framework that links representation learning, generalized planning, temporal control, and resilient coordination into a unified model for event-stream environments.

**Keywords:** Multi-Agent Systems, Artificial Intelligence, Event Stream Management, Generalized Planning, Temporal Reasoning, Stream Processing, Intelligent Scheduling, Resilient Computing, State Representation.

---

### INTRODUCTION

Modern event-driven applications continuously generate streams of observations, transactions, alerts, sensor measurements, and service events. Unlike conventional batch workloads, event streams must often be processed while their arrival patterns, computational requirements, and operational contexts are changing. Consequently, an effective stream-management architecture must simultaneously address event prioritization, routing, resource allocation, workload balancing, failure recovery, and latency-sensitive decision making. A static policy can perform adequately under predictable workloads but may degrade when the system encounters sudden bursts, resource contention, changing event dependencies, or partial component failures.

Artificial intelligence provides a means of moving beyond

static control mechanisms by allowing system components to infer operational states and adapt their decisions. The planning literature is particularly relevant because it treats intelligent control as a process of representing system states, identifying actions, reasoning about their consequences, and constructing policies that achieve desired objectives. Classical planning and search-control research demonstrates the importance of explicitly representing control knowledge rather than treating every decision as an unrestricted search problem (Bacchus, 2001; Baier et al., 2008). Temporal logic further provides a foundation for expressing constraints over sequences of actions and states, which is relevant to event-stream systems in which timing and ordering are fundamental operational properties (Bacchus & Kabanza, 2000).

The multi-agent perspective extends this principle by distributing intelligence across specialized agents. Instead of employing a single centralized controller, an event-stream

architecture can assign distinct responsibilities to observation, prediction, scheduling, routing, resource management, and recovery agents. Such agents can coordinate through shared state abstractions and generalized policies. Recent work on AI-based multi-agent data-streaming models emphasizes the relevance of intelligent coordination to improving resiliency and scalability in event-stream environments (Reddy et al., 2026). The present study develops this direction conceptually by connecting it with generalized planning, symbolic representation, and planning-sketch research.

The primary objective is therefore to formulate an AI-enabled multi-agent framework in which event-stream management is treated as a continuously evolving planning problem. The proposed approach focuses on four research objectives: first, developing compact representations of stream-system states; second, enabling agents to select actions using generalized policies; third, incorporating temporal constraints into stream decisions; and fourth, supporting resilience through decentralized adaptation. The scope is architectural and analytical rather than experimental. The study does not claim measured performance from an implemented system; instead, it derives theoretically grounded findings from the supplied planning literature and applies those principles to event-stream management.

## LITERATURE REVIEW

The supplied literature establishes several theoretical foundations that can be combined to construct an intelligent event-stream management architecture. A first foundation concerns the representation of states and actions. Asai (2019) investigates the grounding of plannable first-order logical representations from images, demonstrating the broader principle that complex environments can be converted into representations suitable for planning. For event-stream management, this idea translates into representing heterogeneous operational observations—such as queue depth, event priority, processor availability, routing status, and failure indicators—as structured states that intelligent agents can reason about.

Description logics provide another important foundation for structured knowledge representation. Baader et al. (2003) establish a formal basis for representing concepts, relationships, and constraints. Within a stream-management system, description-logic-inspired representations can support semantic distinctions between event categories, processing resources, service capabilities, and operational dependencies. The advantage is that agents can reason over structured relationships instead of responding only to raw numerical measurements.

Planning competitions and control mechanisms emphasize the practical importance of search efficiency. Bacchus (2001) demonstrates the relevance of planning systems and their comparative performance, while Baier et al. (2008) investigate procedural control knowledge and preferences beyond classical planning. These contributions are important because an event-stream controller cannot afford unrestricted exploration whenever a new event arrives. It requires mechanisms that constrain decisions according to operational priorities such as latency, resource availability, reliability, and event criticality.

Temporal reasoning is especially relevant to event-stream environments. Bacchus and Kabanza (2000) show how temporal logics can express search-control knowledge for planning. Event streams inherently contain temporal relationships: some events must be processed before others, some must meet deadlines, and some actions are valid only during particular system states. A temporal planning layer can therefore distinguish between merely feasible actions and actions that satisfy timing constraints.

The work of Bonet and Geffner provides a particularly strong theoretical basis for generalized AI policies. Bonet and Geffner (2018) investigate features, projections, and representation change for generalized planning, while Bonet et al. (2019) examine learned features and abstract actions for computing generalized plans. These studies suggest that intelligent planning can operate through abstractions rather than explicitly enumerating every possible state. For event-stream management, abstraction is essential because the number of possible combinations of events, queues, resources, and failures can become extremely large.

Bonet and Geffner (2020a) further address learning first-order symbolic representations for planning from state-space structure. This is directly relevant to adaptive stream-management systems because an agent can potentially derive useful symbolic abstractions from observed operational states. Their work on qualitative numeric planning (Bonet & Geffner, 2020b) also provides a useful conceptual basis for handling numerical resource conditions through qualitative distinctions rather than requiring exact exhaustive enumeration.

General policies and planning width are examined by Bonet and Geffner (2021). This perspective is relevant to scalable event-stream systems because a policy should ideally generalize across multiple operational states rather than being constructed independently for every possible stream condition. Such generalization can reduce the computational burden associated with continuous decision making.

Daggelinckx (2023) considers the generation and verification of planning sketches using temporal logic, while Dalmau-Moreno et al. (2023) investigate task and motion planning through sketch decompositions. These works motivate hierarchical decomposition as a mechanism for controlling

planning complexity. An event-stream management task can similarly be decomposed into stages such as detection, classification, scheduling, execution, and recovery.

Drexler et al. (2021) examine the common subgoal structure of classical planning domains using sketches, and Drexler et al. (2022) extend sketch-based reasoning toward decomposing planning problems into subproblems of bounded width. These ideas are significant for multi-agent systems because individual agents can be assigned bounded subproblems while a coordination mechanism maintains global objectives. Drexler, Francès, and Seipp (2022) also introduce description-logic state features for planning, strengthening the connection between structured state representations and scalable decision making.

Collectively, the literature reveals a gap: the supplied studies strongly develop representations, temporal reasoning, generalized policies, sketches, and planning decomposition, but they do not directly formulate these principles as an integrated architecture for AI-driven event-stream management. The proposed framework addresses this conceptual gap by treating stream management as a distributed planning environment in which agents operate over structured states, generalized policies, temporal constraints, and decomposed objectives.

## METHODOLOGY

### 3.1 Conceptual Architecture

The proposed architecture consists of six interacting functional layers: event observation, state abstraction, intelligent planning, multi-agent coordination, execution, and resilience management. The architecture is designed around the assumption that an event stream is not merely a sequence of messages but a dynamic state-transition environment.

The event-observation layer receives incoming events and extracts operational attributes such as event type, priority, timestamp, dependency, processing requirement, and urgency. The observation process transforms heterogeneous stream data into a structured state representation. Inspired by symbolic grounding and structured planning representations, the objective is to ensure that downstream agents reason over meaningful operational states rather than unstructured event payloads (Asai, 2019).

The state-abstraction layer converts detailed observations into compact features. For example, instead of representing every individual queue measurement, the system may classify queue conditions as low, moderate, high, or critical. Similarly, resources can be represented

through qualitative availability states. This abstraction follows the broader rationale of generalized planning and qualitative reasoning, where compact features support policies across multiple states (Bonet & Geffner, 2018; Bonet & Geffner, 2020b).

### 3.2 Multi-Agent Functional Decomposition

The framework employs specialized agents with complementary responsibilities. A Stream Monitoring Agent observes event rates, queue conditions, latency, and failures. A Classification Agent determines event priority and processing requirements. A Resource Planning Agent evaluates available computational resources and determines suitable execution locations. A Routing Agent selects processing paths based on workload and service conditions. A Recovery Agent responds to failures by identifying alternative execution strategies. Finally, a Coordination Agent maintains consistency between local decisions and global system objectives.

This decomposition follows the principle that complex planning problems can be divided into structured subproblems. Planning sketches and bounded-width decomposition provide theoretical support for this approach because they reduce large planning problems into manageable components (Drexler et al., 2021; Drexler et al., 2022).

### 3.3 Generalized Policy Mechanism

Rather than generating a new plan for every event, the architecture uses generalized policies. A policy can specify conditional actions such as: if event priority is critical and processing capacity is available, assign the event to the lowest-latency resource; if capacity is insufficient, invoke an alternative resource; if the selected resource fails, activate a recovery route.

This design reflects the importance of generalized policies and abstract actions identified by Bonet et al. (2019) and Bonet and Geffner (2021). The benefit is computational scalability. Instead of treating every combination of event and infrastructure state as unique, the agent applies an abstract policy to a class of states.

### 3.4 Temporal Decision Layer

Event-stream systems require decisions that consider time explicitly. The proposed temporal layer maintains constraints involving deadlines, ordering, processing windows, and event dependencies. Temporal logic provides the theoretical basis for expressing these constraints (Bacchus & Kabanza, 2000).

For example, a high-priority alert may require processing before a defined deadline. A dependent event may not be executed until its predecessor has completed. If the primary processing resource becomes unavailable, the system should select an alternative action without violating the temporal constraint. Planning sketches verified with temporal logic

provide a conceptual basis for enforcing such conditions (Daggelinckx, 2023).

### 3.5 Coordination and Conflict Resolution

Decentralized agents can produce locally rational decisions that conflict globally. For example, a routing agent may select a high-performance resource while a resource-management agent simultaneously assigns that resource to another critical workload. The coordination layer therefore maintains a shared abstract state and applies priority rules.

Description-logic state features can support this shared representation by expressing relationships among events, resources, capabilities, and dependencies (Drexler, Francès, & Seipp, 2022). The coordination mechanism can consequently distinguish between compatible and conflicting actions before execution.

### 3.6 Resilience Mechanism

Reliability is achieved through continuous monitoring, alternative policies, and recovery planning. If an execution node fails, the Recovery Agent identifies a replacement path using the current abstract state. Rather than reconstructing the complete planning problem, the system uses a bounded recovery subproblem. This is consistent with sketch-based decomposition approaches that seek to control planning complexity (Drexler et al., 2022).

The proposed approach is also aligned with the broader objective of AI-based multi-agent event-stream models in improving resiliency and scalability (Reddy et al., 2026). The essential distinction is that resilience is treated not simply as redundancy but as an intelligent decision process involving state recognition, policy selection, and adaptive recovery.

### 3.7 Analytical Evaluation Criteria

Because the present work is a conceptual research framework rather than an empirical implementation, evaluation is defined through architectural criteria rather than fabricated numerical measurements. The principal criteria are decision latency, policy generalization, resource utilization, fault-recovery responsiveness, coordination overhead, and scalability with respect to the number of events and agents.

A successful implementation would be expected to demonstrate that abstraction reduces planning complexity, generalized policies reduce repeated computation, temporal reasoning improves deadline compliance, and agent decomposition limits the computational burden imposed on individual decision modules.

## RESULTS

The analytical findings indicate that the proposed architecture provides a coherent mechanism for transforming event-stream management from a predominantly reactive process into an adaptive planning process. The first major finding concerns state representation. Raw stream data contain substantial operational detail but are not necessarily suitable for direct decision making. Structured and abstract representations provide a more useful intermediate layer because agents can reason about classes of states rather than individual observations. This finding is consistent with the role of representation change, learned features, and symbolic state structures in generalized planning (Asai, 2019; Bonet & Geffner, 2018; Drexler, Francès, & Seipp, 2022).

The second finding concerns policy generalization. A high-volume event-stream system may encounter an effectively unbounded number of event/resource combinations. Constructing a unique plan for every state would therefore be computationally undesirable. Generalized policies and abstract actions offer a more scalable alternative because one policy can govern multiple operational states (Bonet et al., 2019; Bonet & Geffner, 2021). In practical terms, an agent can learn or construct policies around conditions such as congestion, priority, resource failure, and deadline pressure rather than around individual event identities.

The third finding is that temporal reasoning is fundamental rather than optional. Event streams contain ordering and deadline constraints that cannot be represented adequately through static resource optimization alone. A system may possess sufficient processing capacity yet still fail if the selected action violates an event dependency or deadline. Temporal control therefore provides a decision layer capable of distinguishing between technically executable and operationally valid actions (Bacchus & Kabanza, 2000).

The fourth finding concerns decomposition and resilience. Dividing global stream-management objectives among specialized agents can reduce local planning complexity. Sketch-based decomposition provides theoretical support for this approach by constraining the search space and allowing complex planning objectives to be divided into manageable subproblems (Dalmau-Moreno et al., 2023; Drexler et al., 2021; Drexler et al., 2022). In a failure scenario, the Recovery Agent can focus on the bounded problem of restoring an affected processing path instead of recomputing the entire system policy.

The fifth finding is that multi-agent coordination introduces a new trade-off. Decentralization improves modularity and potentially increases scalability, but coordination itself generates communication and consistency overhead. Consequently, the optimal architecture should not maximize the number of agents indiscriminately. Instead, agents should be organized around functionally meaningful planning

boundaries.

Overall, the analytical findings support the central proposition that combining symbolic representations, generalized policies, temporal reasoning, and decomposed multi-agent control provides a theoretically stronger foundation for resilient event-stream management than isolated reactive rules. This direction is consistent with the objective of AI-based multi-agent models for improving resiliency and scalability in event streaming (Reddy et al., 2026).

## DISCUSSION

The principal theoretical implication of the proposed framework is that event-stream management can be interpreted as a generalized planning problem. Conventional stream-processing approaches frequently emphasize throughput and latency, but intelligent management additionally requires reasoning about state transitions and action consequences. The planning literature demonstrates that control knowledge, representations, and generalized policies can substantially influence the efficiency of intelligent decision processes (Bacchus, 2001; Baier et al., 2008). Applying these principles to event streams creates a conceptual bridge between AI planning and distributed stream infrastructure.

A major practical implication is the potential reduction of computational decision complexity. If each incoming event requires independent optimization, control overhead can become significant. Generalized policies allow an agent to reuse decision structures across related states (Bonet & Geffner, 2021). However, policy reuse introduces a risk: a policy that generalizes too aggressively may overlook rare operational conditions. The architecture therefore requires a balance between abstraction and state-specific reasoning.

Temporal reasoning presents a second important advantage. Event-stream management frequently involves deadlines and dependencies, meaning that resource allocation cannot be optimized independently from event order. The temporal-logical perspective of Bacchus and Kabanza (2000) and the temporal verification orientation of Daggelinckx (2023) support the incorporation of explicit temporal constraints. Nevertheless, temporal reasoning can increase model complexity when event dependencies become highly interconnected.

The multi-agent design also produces both benefits and limitations. Specialized agents improve modularity and permit local decision making, but distributed decisions may conflict. A coordination layer is therefore essential. The use of structured state features and planning sketches can reduce such conflicts by ensuring that agents operate

over common abstractions (Drexler, Francès, & Seipp, 2022). Yet centralized coordination can itself become a bottleneck, creating a tension between global consistency and decentralization.

Another limitation is that the proposed framework has not been validated through a physical or simulated implementation in this study. Accordingly, claims concerning improved throughput, latency, or fault-recovery time should be regarded as hypotheses for future experimental verification rather than established empirical outcomes. This distinction is important for research validity. The architecture establishes theoretically motivated mechanisms, but implementation-specific factors—including communication overhead, policy-learning cost, infrastructure heterogeneity, and workload distributions—will determine actual performance.

The framework also extends the perspective represented by AI-based multi-agent event-stream research by connecting resilience with planning and representation. The objective of improving streaming resiliency and scalability through multi-agent intelligence is directly aligned with Reddy et al. (2026). The present framework, however, emphasizes why such systems may benefit from generalized planning, temporal constraints, and sketch-based decomposition rather than relying exclusively on reactive intelligence.

Finally, the literature suggests that representation quality is likely to be a decisive factor. Poor abstractions can conceal operationally important distinctions, whereas overly detailed representations can produce computationally expensive planning spaces. The balance between these extremes is central to future implementation. Learned first-order representations, qualitative numerical reasoning, and description-logic state features offer complementary theoretical mechanisms for investigating this balance (Bonet & Geffner, 2020a; Bonet & Geffner, 2020b; Drexler, Francès, & Seipp, 2022).

## CONCLUSION

This paper presented a conceptual AI-enabled multi-agent approach for reliable and high-performance event-stream management. The framework integrates structured state representation, generalized planning policies, temporal reasoning, multi-agent decomposition, coordination, and intelligent recovery. Its theoretical foundation is derived from the supplied literature on planning, symbolic representation, generalized policies, temporal logic, and planning sketches.

The analysis indicates that event-stream management can be formulated as a dynamic planning environment in which agents continuously interpret operational states and select actions according to generalized policies and temporal constraints. Representation abstraction can reduce decision complexity, while policy generalization can support scalable control across multiple operational states. Temporal reasoning

addresses deadlines and event dependencies, and sketch-based decomposition provides a mechanism for managing complex recovery and coordination tasks.

The principal contribution is therefore an integrated research framework rather than a claim of experimentally measured performance. Future work should implement the architecture in a distributed stream-processing environment and evaluate it using controlled workloads, burst conditions, resource failures, heterogeneous event priorities, and varying numbers of agents. Quantitative evaluation should examine throughput, end-to-end latency, recovery time, coordination overhead, resource utilization, and policy generalization. Such experiments would determine whether the theoretical advantages identified in the planning literature translate into measurable improvements in real-world event-stream systems.

## REFERENCES

1. Asai, M. (2019). Unsupervised grounding of plannable first-order logic representation from images. In Lipovetzky, N., Onaindia, E., & Smith, D. E. (Eds.), *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling (ICAPS 2019)*, pp. 583–591. AAAI Press.
2. Baader, F., Calvanese, D., McGuinness, D. L., Nardi, D., & Patel-Schneider, P. F. (Eds.). (2003). *The Description Logic Handbook: Theory, Implementation and Applications*. Cambridge University Press.
3. Bacchus, F. (2001). The AIPS'00 planning competition. *AI Magazine*, 22(3), 47–56.
4. Bacchus, F., & Kabanza, F. (2000). Using temporal logics to express search control knowledge for planning. *Artificial Intelligence*, 116(1–2), 123–191.
5. Baier, J. A., Fritz, C., Bienvenu, M., & McIlraith, S. A. (2008). Beyond classical planning: Procedural control knowledge and preferences in state-of-the-art planners. In *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence (AAAI 2008)*, pp. 1509–1512. AAAI Press.
6. Bonet, B., Francès, G., & Geffner, H. (2019). Learning features and abstract actions for computing generalized plans. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence (AAAI 2019)*, pp. 2703–2710. AAAI Press.
7. Bonet, B., & Geffner, H. (2018). Features, projections, and representation change for generalized planning. In Lang, J. (Ed.), *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI 2018)*, pp. 4667–4673. IJCAI.
8. Bonet, B., & Geffner, H. (2020a). Learning first-order symbolic representations for planning from the structure of the state space. In De Giacomo, G. (Ed.), *Proceedings of the 24th European Conference on Artificial Intelligence (ECAI 2020)*, pp. 2322–2329. IOS Press.
9. Bonet, B., & Geffner, H. (2020b). Qualitative numeric planning: Reductions and complexity. *Journal of Artificial Intelligence Research*, 69, 923–961.
10. Bonet, B., & Geffner, H. (2021). General policies, representations, and planning width. In Leyton-Brown, K., & Mausam (Eds.), *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*, pp. 11764–11773. AAAI Press.
11. Daggelinckx, A. (2023). Generating and verifying planning sketches using temporal logic. Master's thesis, Utrecht University.
12. Dalmau-Moreno, M., García, N., Gómez, V., & Geffner, H. (2023). Combined task and motion planning via sketch decompositions. In *ICAPS 2023 Planning and Robotics Workshop (Plan-Rob)*.
13. Drexler, D., Francès, G., & Seipp, J. (2022). Description logics state features for planning (DLPlan).
14. Drexler, D., Seipp, J., & Geffner, H. (2021). Expressing and exploiting the common subgoal structure of classical planning domains using sketches. In Erdem, E., Bienvenu, M., & Lakemeyer, G. (Eds.), *Proceedings of the Eighteenth International Conference on Principles of Knowledge Representation and Reasoning (KR 2021)*, pp. 258–268. IJCAI Organization.
15. Drexler, D., Seipp, J., & Geffner, H. (2022). Learning sketches for decomposing planning problems into subproblems of bounded width. In Thiébaux, S., & Yeoh, W. (Eds.), *Proceedings of the Thirty-Second International Conference on Automated Planning and Scheduling (ICAPS 2022)*, pp. 62–70. AAAI Press.
16. R. Reddy, P. Udayaraju, K. K. Goyal, S. C. R. Vudem, R. Sayana and V. Gummati, "Creating an AI-based Multi-Agent Model for Optimized Data Streaming with Improved Resiliency and Scalability for Event Streaming," 2026 6th International Conference on Image Processing and Capsule Networks (ICIPCN), Dhulikhel, Nepal, 2026, pp. 1372-1379, doi: 10.1109/ICIPCN67432.2026.11438445.