Frontiers in Emerging Artificial Intelligence and Machine Learning (Open Access)

Volume 03, Issue 08, August 2026,
Publish Date: 22-08-2026
DOI: https://doi.org/10.64917/feaiml/Volume03Issue08-01 **PageNo.01-06**

Intelligent Multi-Agent Generative AI Framework for Autonomous Exception Resolution in SAP S/4HANA Manufacturing

Arjun Mehta
AI Research and Innovation Division, Intelligent Systems Research Center, India

ABSTRACT

Manufacturing environments operating on SAP S/4HANA increasingly depend on real-time coordination among production planning, inventory, procurement, maintenance, quality, logistics, and order-management processes. Exception conditions such as material shortages, production delays, capacity conflicts, quality deviations, and delivery disruptions therefore require decisions that are both rapid and context-aware. This paper proposes an Intelligent Multi-Agent Generative AI Framework for Autonomous Exception Resolution in SAP S/4HANA Manufacturing. The framework conceptualizes manufacturing exception management as a distributed decision problem in which specialized AI agents detect, interpret, negotiate, validate, and resolve operational deviations while maintaining authorization and security controls. The theoretical foundation is derived exclusively from the supplied literature, particularly research on delegation, proxy signatures, secure mobile agents, signcryption, identity-based security, and publicly verifiable cryptographic mechanisms. These studies provide a conceptual basis for trusted delegation, authenticated agent interaction, confidentiality, integrity, and verifiable autonomous actions (Mambo et al., 1996; Kim et al., 1997; Lee et al., 2001; Chow et al., 2004). The proposed framework integrates event interpretation, multi-agent reasoning, policy-constrained action selection, cryptographically protected agent communication, and human escalation. Analytical findings indicate that autonomous exception resolution should not be implemented as unrestricted generative decision-making; instead, it requires bounded delegation, explicit authorization, verifiable actions, and exception-specific governance. The resulting architecture provides a research-oriented model for connecting generative AI reasoning with enterprise manufacturing execution while addressing the trust and security challenges inherent in autonomous operational decisions.

KEYWORDS: Generative AI, Multi-Agent Systems, SAP S/4HANA, Manufacturing, Autonomous Exception Resolution, Intelligent Agents, Delegated Authorization, Signcryption, Enterprise AI.

INTRODUCTION

1.1 Background

Modern manufacturing is characterized by interconnected operational dependencies. A deviation in one process can propagate rapidly into several others. A shortage of a critical component, for example, may affect production orders, material requirements, procurement decisions, delivery commitments, and customer service simultaneously. Conventional exception management frequently relies on predefined rules, manual intervention, or isolated application workflows. Such mechanisms are effective for predictable conditions but become less effective when multiple constraints interact and the appropriate response requires contextual reasoning.

Generative artificial intelligence introduces the possibility of interpreting complex operational information and generating candidate responses rather than merely executing predetermined rules. A manufacturing-oriented generative AI system can potentially summarize an exception, identify its causal context, compare alternative responses, and formulate an action plan. However, a single general-purpose model is not necessarily appropriate for executing cross-functional enterprise decisions. Manufacturing exceptions frequently require specialized knowledge spanning production, inventory, procurement, quality, maintenance, and logistics.

A multi-agent architecture provides an alternative. Different intelligent agents can be assigned specialized responsibilities while communicating through controlled interfaces. This approach also creates a natural foundation for delegated authority. The historical literature on proxy

signatures demonstrates how signing authority can be delegated from one entity to another under defined conditions (Mambo et al., 1996). Subsequent research on proxy signatures and their applications further emphasizes the importance of controlled delegation in distributed environments (Kim et al., 1997). These principles are relevant to autonomous enterprise AI because an AI agent acting on behalf of an authorized human or organizational role requires more than computational capability; it requires demonstrable authority.

The proposed Intelligent Multi-Agent Generative AI Framework for Autonomous Exception Resolution in SAP S/4HANA Manufacturing therefore treats autonomy as a controlled delegation problem rather than unrestricted automation.

1.2 Problem Statement

The principal research problem is how an AI system can autonomously resolve manufacturing exceptions without compromising decision accountability, authorization, confidentiality, or operational consistency. Generative AI can produce plausible recommendations, but plausibility does not constitute authorization. An agent may generate an appropriate production adjustment but still lack permission to modify a production order, release a purchase requirement, alter a delivery commitment, or initiate another consequential transaction.

A second problem concerns coordination. A manufacturing exception may require multiple specialized decisions. An inventory agent may identify a shortage, a production agent may determine its effect on manufacturing orders, a procurement agent may evaluate replenishment, and a logistics agent may assess delivery implications. Independent agents therefore require secure communication and a mechanism for validating the collective decision.

1.3 Research Relevance

The supplied cryptographic literature provides a useful theoretical foundation for this problem. Diffie and Heliman established foundational principles for secure communication through public-key cryptography (Diffie & Heliman, 1976). Signcryption subsequently combined confidentiality and authentication objectives within a unified cryptographic operation (Zheng, 1997; Lee & Mao, 2003). Identity-based signcryption extended these ideas toward identity-oriented security architectures (Malone-Lee, 2002; Libert & Quisquater, 2003). These concepts can be translated conceptually into requirements for secure AI-agent communication.

The framework also draws upon the delegation concepts represented by proxy-signature research (Mambo et al., 1996; Kim et al., 1997). In an enterprise AI setting, delegated authority can provide a mechanism through which an AI agent performs an action within explicitly defined organizational boundaries.

1.4 Objectives, Scope, and Significance

The objective is to develop a conceptual architecture for autonomous manufacturing exception resolution that combines multi-agent reasoning, generative AI, delegated authorization, secure communication, action verification, and human escalation.

The scope is limited to manufacturing exception management within an SAP S/4HANA-oriented enterprise environment. The paper does not claim empirical validation of production performance or actual SAP transaction execution. Instead, it proposes a research framework and derives analytical findings from the theoretical relationships among autonomous agents, delegation, security, and exception resolution.

The significance of the framework lies in repositioning enterprise generative AI from a recommendation-only technology toward a governed autonomous decision mechanism. Its central proposition is that useful autonomy requires both reasoning capability and institutional trust.

LITERATURE REVIEW

2.1 Cryptographic Foundations for Trusted Agent Interaction

The foundations of secure digital communication can be traced to public-key cryptographic concepts presented by Diffie and Heliman (1976). Their work establishes the conceptual basis for communicating securely in environments where participants cannot simply rely on a shared secret established through a trusted physical channel. For autonomous enterprise agents, this principle is important because distributed agents must exchange information without assuming unrestricted trust.

Zheng (1997) introduced signcryption as an approach intended to combine signature and encryption functions while reducing the conceptual cost of performing the two operations independently. Lee and Mao (2003) further examined signcryption using RSA. In an AI-agent architecture, the underlying principle is relevant because agent messages may require both confidentiality and evidence of origin. An exception-management agent should not merely receive a recommendation; it should be possible to determine whether the message originated from an authorized computational participant and whether its contents were protected against unauthorized disclosure or manipulation.

Identity-based signcryption extends this security perspective. Malone-Lee (2002) examined identity-based signcryption, while Libert and Quisquater (2003) proposed an identity-based signcryption scheme derived from pairings. Although these works are cryptographic rather than enterprise-AI studies, they provide a theoretical basis for designing identity-aware agent communication.

2.2 Delegation and Autonomous Agents

Proxy signatures provide an especially relevant conceptual connection to autonomous AI. Mambo et al. (1996) investigated delegation of signing power, demonstrating a security model in which one party can authorize another party to perform signing operations. Kim et al. (1997) revisited proxy signatures and examined their role in distributed security. These ideas map naturally onto enterprise AI delegation: an AI agent may act for a planner, supervisor, or organizational function, but only within a specified authorization boundary.

The work by Lee, Kim, and Kim on strong non-designated proxy signatures (2001) further highlights the importance of controlling delegated authority in mobile-agent environments. Their research suggests that delegation must be treated as a security property rather than merely a convenience mechanism. This distinction is critical for manufacturing AI because an agent capable of modifying operational data can create significant downstream consequences.

The subsequent study of strong proxy signatures and their applications by Lee, Kim, and Kim (2001) reinforces the relationship between delegated action and application-level security. In the proposed framework, agent autonomy consequently depends on a policy layer that determines which actions an agent may execute, under what conditions, and with what level of authorization.

2.3 Verifiability and Distributed Coordination

Chow et al. (2004) addressed public verifiability and public ciphertext authenticity within an identity-based signcryption scheme. These concepts are particularly relevant to autonomous multi-agent systems because a distributed decision should ideally be auditable after execution. If multiple agents contribute to an exception-resolution process, the enterprise should be able to reconstruct the decision chain and determine whether messages and actions were authentic.

The supplied literature therefore collectively suggests four foundational properties for autonomous enterprise agents: secure communication, identity, controlled delegation, and verifiability. What is missing from the supplied literature is an integrated manufacturing-specific architecture that combines these properties with generative AI reasoning. The proposed framework addresses this conceptual gap.

2.4 Research Gap and Theoretical Positioning

The principal gap is not the absence of individual security mechanisms. Rather, it is the absence of a unified conceptual model linking secure delegated authority with generative multi-agent reasoning for manufacturing exception resolution. Existing cryptographic studies primarily establish security mechanisms, while the proposed research applies their underlying principles to autonomous enterprise decision processes.

The theoretical position of this study is therefore that autonomous AI action should be modeled as cryptographically accountable delegation. Generative reasoning produces candidate decisions, while authorization, security, policy validation, and verification determine whether those decisions can become executable enterprise actions.

3. METHODOLOGY

3.1 Research Design

This study adopts a conceptual design-science-oriented methodology. The objective is to construct a logically coherent framework rather than report experimental performance. The methodology involves four stages: identification of exception-management requirements, decomposition into specialized agents, integration of delegated security controls, and formulation of an autonomous resolution lifecycle.

The framework is organized around the principle that an exception passes through five logical states: detection, interpretation, coordination, authorization, and execution/verification.

3.2 Proposed Multi-Agent Architecture

The architecture contains a central orchestration layer and specialized functional agents. The Exception Detection Agent identifies abnormal operational conditions from manufacturing events. The Context Analysis Agent consolidates relevant production, material, capacity, and order information. The Production Agent evaluates manufacturing implications, while the Inventory Agent analyzes material availability. The Procurement Agent evaluates replenishment alternatives, and the Logistics Agent examines downstream delivery effects.

A Generative Reasoning Agent synthesizes the outputs of these specialized agents. Rather than independently inventing operational actions, it generates candidate resolution strategies from verified contextual information. An Authorization Agent then determines whether a

proposed action falls within the delegated authority assigned to the requesting agent.

This structure creates separation between reasoning and execution. Such separation is important because a generative model may produce a technically plausible recommendation without possessing authority to execute it.

3.3 Exception Resolution Lifecycle

The lifecycle begins when an exception event is detected. Consider a hypothetical shortage of a component required for an imminent production order. The Exception Detection Agent classifies the shortage and forwards it to the Context Analysis Agent.

The Context Analysis Agent determines affected production orders, available inventory, expected replenishment, and operational dependencies. Specialized agents then independently evaluate alternatives. The Production Agent might recommend rescheduling, while the Procurement Agent could recommend expedited replenishment.

The Generative Reasoning Agent synthesizes these alternatives and generates a ranked resolution plan. The plan is subsequently subjected to policy validation. If the required action falls within delegated authority, it can proceed toward execution. If it exceeds the permitted authority threshold, the case is escalated to a human decision-maker.

3.4 Delegated Authorization Model

Delegation is central to the proposed framework. The conceptual authorization structure is inspired by proxy-signature research, where authority can be delegated from an original signer to another entity (Mambo et al., 1996; Kim et al., 1997).

For manufacturing AI, delegation can be represented as:

Principal → Delegation Policy → AI Agent → Authorized Action

A policy may define transaction type, monetary or operational threshold, organizational scope, validity period, affected manufacturing plant, and escalation conditions. For example, an inventory agent could be authorized to recommend stock reallocation but prohibited from changing safety-stock parameters. A procurement agent could initiate a replenishment proposal but require human approval for high-value purchases.

This model prevents the assumption that intelligence automatically implies authority.

3.5 Secure Agent Communication

Agent-to-agent communication should provide confidentiality, authentication, integrity, and evidence of origin. The conceptual security layer therefore incorporates principles derived from public-key cryptography (Diffie & Heliman, 1976), signcryption (Zheng, 1997), RSA-based signcryption (Lee & Mao, 2003), and identity-based signcryption (Malone-Lee, 2002; Libert & Quisquator, 2003).

A message exchanged between two agents can conceptually contain an event identifier, sender identity, timestamp, contextual data, proposed action, authorization scope, and cryptographic protection. The receiving agent can then verify that the information originates from an expected source and has not been altered.

This architecture is especially important where multiple agents collaborate on a consequential manufacturing decision.

3.6 Verification and Auditability

The proposed framework incorporates a verification layer inspired by public-verifiability principles discussed by Chow et al. (2004). Every autonomous action should produce an auditable decision record containing the triggering exception, contextual evidence, participating agents, generated alternatives, selected action, authorization status, and execution outcome.

The objective is not simply to log AI activity but to create a traceable relationship between exception, reasoning, delegation, and action. Such traceability supports accountability and enables post-event analysis.

3.7 Human Escalation

Autonomy should be adaptive rather than absolute. Low-risk exceptions can potentially be resolved automatically when the action remains inside a predefined authorization envelope. High-impact or ambiguous cases should be escalated.

A practical decision model can therefore classify actions into three categories: automatically executable, conditionally executable, and human-authorized. This structure limits the operational impact of generative-model uncertainty while retaining the efficiency benefits of automation.

4. RESULTS

The conceptual analysis produces five principal findings. First, autonomous exception resolution is most appropriately modeled as a multi-stage decision process rather than a single generative prediction task. Manufacturing exceptions often involve several interacting constraints. Separating detection, contextual interpretation, domain-specific analysis, reasoning, authorization, and execution reduces the risk that one model will

independently make an unsupported enterprise decision. The multi-agent structure therefore improves functional specialization and creates explicit control points.

Second, delegation emerges as a fundamental requirement for enterprise AI autonomy. The proxy-signature literature demonstrates that authority can be transferred under defined conditions rather than assumed implicitly (Mambo et al., 1996; Kim et al., 1997). Applying this principle to manufacturing suggests that AI agents should receive explicit operational permissions. An agent should not gain authority simply because a generative model is capable of producing a particular transaction or recommendation.

Third, security must be integrated with reasoning rather than added after decision generation. The literature on public-key cryptography and signcryption establishes the importance of secure identity, confidentiality, authentication, and integrity (Diffie & Heliman, 1976; Zheng, 1997). For a multi-agent manufacturing environment, these properties apply to both information exchange and action authorization. A secure communication channel alone is insufficient if the receiving agent cannot determine whether the sender is authorized to request the proposed operation.

Fourth, verifiability is necessary for responsible autonomy. The public-verifiability perspective in identity-based signcryption provides a useful theoretical foundation for constructing auditable AI-agent interactions (Chow et al., 2004). In manufacturing, an autonomous action should be reconstructable from its triggering exception and participating decision components. This requirement transforms AI from an opaque recommendation mechanism into an accountable enterprise process.

Fifth, the framework indicates that autonomy should be risk-sensitive. Routine, reversible actions can be assigned broader autonomous authority, while irreversible, financially significant, safety-sensitive, or strategically important actions should require additional validation. The resulting architecture balances operational speed against governance. Its principal limitation is that these findings remain conceptual: the framework has not been empirically tested against actual SAP S/4HANA transaction workloads, latency measurements, exception-resolution accuracy, or production environments.

5. DISCUSSION

The findings suggest that the central challenge in enterprise generative AI is not simply improving model intelligence but determining how intelligence can be converted into trustworthy action. This distinction is particularly important in manufacturing, where a seemingly minor exception can propagate through interconnected processes. A generative model may identify an attractive solution, but an autonomous enterprise system must additionally establish whether the solution is supported by reliable context, permitted by organizational policy, and safe to execute.

The delegation model provides an important theoretical mechanism for addressing this problem. Proxy signatures demonstrate that authority can be delegated without eliminating the distinction between the original authority holder and the delegated actor (Mambo et al., 1996; Kim et al., 1997). In the proposed architecture, this distinction maps onto the relationship between organizational decision-makers and AI agents. An agent can act autonomously while remaining institutionally subordinate to explicit authorization policies. This approach is preferable to unrestricted agent autonomy because it establishes boundaries before execution.

The cryptographic literature also creates an important conceptual connection between communication security and AI accountability. Signcryption combines security objectives that are particularly relevant to distributed autonomous systems (Zheng, 1997; Lee & Mao, 2003). Identity-based approaches further demonstrate how identity can be incorporated into cryptographic protection (Malone-Lee, 2002; Libert & Quisquator, 2003). In a manufacturing multi-agent environment, these principles suggest that every consequential interaction should be associated with a verifiable computational identity.

A further implication concerns explainability. Generative AI explanations are not equivalent to auditability. A natural-language explanation may describe why an agent selected an action, but an enterprise audit requires evidence concerning which exception triggered the process, what information was considered, which agents participated, what authorization was applicable, and what action was executed. The proposed event-and-authorization record therefore provides a stronger accountability mechanism than narrative explanation alone.

There are also important trade-offs. Greater autonomy can reduce response time but increase the consequences of erroneous reasoning. More extensive verification can improve trust but introduce processing overhead. Specialized agents improve functional separation but increase orchestration complexity. Strict authorization limits reduce operational risk but can prevent useful automation in unusual circumstances. Consequently, the appropriate objective is not maximum autonomy but controlled autonomy optimized for operational risk.

The framework has several limitations. The supplied literature primarily concerns cryptographic and delegation mechanisms rather than generative AI or SAP manufacturing processes. Therefore, the connection between these domains is theoretically derived rather than empirically

demonstrated. The study also does not evaluate model hallucination rates, transaction-processing latency, cybersecurity attack resistance, or production-scale agent coordination. Future research should validate the framework using simulated manufacturing exceptions and, subsequently, controlled enterprise environments.

6. CONCLUSION

This paper proposed an Intelligent Multi-Agent Generative AI Framework for Autonomous Exception Resolution in SAP S/4HANA Manufacturing. The framework conceptualizes exception management as a controlled sequence of detection, contextual reasoning, multi-agent coordination, delegated authorization, secure execution, and verification.

The analysis demonstrates that autonomous enterprise AI requires more than generative reasoning. Delegation principles from proxy-signature research provide a theoretical basis for bounded agent authority, while public-key cryptography, signcryption, identity-based security, and verifiability provide foundations for trusted agent communication and accountability (Diffie & Heliman, 1976; Mambo et al., 1996; Zheng, 1997; Chow et al., 2004). The proposed architecture consequently separates intelligence from authority: an agent may generate a solution, but a policy-controlled mechanism determines whether that solution can become an executable enterprise action.

The principal contribution is therefore a governance-oriented conception of multi-agent generative AI in manufacturing. Future research should operationalize the framework through simulation, benchmark exception-resolution accuracy and latency, evaluate agent coordination under conflicting recommendations, and investigate formal authorization mechanisms for production-scale deployment. Such studies can determine whether the proposed model can achieve measurable improvements in manufacturing responsiveness while preserving security, traceability, and organizational control.

REFERENCES

1. Chow SSM, Yiu SM, Hui LCK, Chow KP. Efficient forward and provably secure ID based signcryption scheme with public verifiability and public ciphertext authenticity. In Procceedings of Information Security and Cryptology (ICISC'03), Lim JI, Lee DH (eds), LNCS 2971. Springer-Verlag: Seoul, Korea, 2004; 352–369.
2. Diffie W, Heliman M. New direction in cryptography. IEEE Transactions on Information Theory 1976; 22 (6): 644-654.
3. Kim S, Park S, and Won D. Proxy signatures, revisited. In Proceedings of ICICS 97, LNCS 1334, Springer-Verlag; 1997; 223-232.
4. Kim S, Park S, Won D. Proxy signatures, revisited. In Proceedings of Information and Communications Security. (ICICS' 97), Han Y, Okamoto T, Qing S (eds), LNCS 1334. Springer- Verlag:Beijing, China, 1997; 223-232.
5. Lee B, Kim H, Kim K. Secure mobile agent using strong non-designated proxy signature. In Proceedings of Information Security and Privacy (ACISP'01), Varadharajan V, Mu Y (eds), LNCS 2119. Springer Verlag: Sydney, Australia, 2001; 474-486.
6. Lee B, Kim H, Kim K. Strong proxy signature and its applications, In Proceedings of the Symposium on Cryptography and Information Security (SCIS'01), Oiso, Japan,2001;603-608.
7. Lee JM, Mao W. Two birds one stone: signcryption using RSA. In Topics in Cryptology (CT-RSA'03), Joye M (ed), LNCS 2612. Springer-Verlag: San Francisco, CA, USA, 2003; 211–225.
8. Libert B, Quisquator JJ. A new identity based signcryption scheme from pairings. In Proceedings of IEEE Information Theory Workshop (ITW'03). Elsevier: Paris, France, 2003; 155–158.
9. Mambo M, Usuda K, Okamoto E. Proxy signature: delegation of the power to sign messages. IEICE Transaction on Fundamentals 1996; E79-A (9): 1338-1353.
10. Malone-Lee J. Identity based signcryption. Available from: http://eprint.iacr.org/2002/098.pdf [Accessed on 30 May 2011].
11. Zheng Y. Digital Signcryption or How to Achieve Cost (Signature & Encryption) << Cost (Signature) + Cost (Encryption)", Advances in Cryptology – CRYPTO' 97, volume 1294 of Lecture Notes in Computer Science, Springer-Verlag; 1997; 165-179.
12. Kalal, M. (2025). Autonomous Exception Management in SAP S/4HANA Manufacturing Through Multi-agent Generative AI and Event-driven Supply Networks.