Третьякова Вероника Сергеевна
НИУ ВШЭ
Tretyakova Veronika Sergeevna
HSE University

AGENT ORCHESTRATION AS A FOUNDATION FOR INTELLIGENT BUSINESS PROCESS MANAGEMENT IN THE DIGITAL ENTERPRISE ENVIRONMENT

Аннотация. В статье рассматривается агентная оркестрация в интеллектуальном управлении бизнес-процессами. Предложена архитектура, включающая прием заданий, координацию, исполнение агентами, память, интеграцию и управление. Показано, что эффективность зависит от формализации задач, контроля доступа, проверки результатов, журналирования и участия человека. Мультиагентные системы могут сокращать длительность процессов и повышать качество решений при работе в контролируемой среде.

Abstract. The article examines agent orchestration in intelligent business process management. A functional architecture covering task intake, coordination, agent execution, memory, integration, and governance is proposed. It is shown that effectiveness depends on task formalization, access control, validation, auditability, and human approval. Multi-agent systems can reduce process duration and improve decision support when integrated into controlled workflows.

Ключевые слова: Агентная оркестрация, мультиагентные системы, управление бизнес-процессами, интеллектуальная автоматизация, no-code платформы, поддержка решений, управление ИИ.

Keywords: Agent orchestration, multi-agent systems, business process management, intelligent automation, no-code platforms, decision support, AI governance.

Introduction

Enterprise use of artificial intelligence (AI) is moving from isolated analytical tools toward systems capable of planning and executing multi-step work. In 2025, 88% of surveyed organizations reported regular AI use in at least one business function, while 62% were experimenting with AI agents. Nearly two-thirds had not yet begun enterprise-wide scaling, which indicates that access to models is advancing faster than process redesign and operational control [1]. Traditional business process management (BPM) and robotic process automation provide stable execution when rules, inputs, and system responses are predictable. Their effectiveness declines in processes involving unstructured documents, changing context, several decision points, and frequent exceptions. Autonomous agents can interpret information and select actions, yet isolated deployment creates duplicated work, inconsistent state, excessive tool access, and weak accountability.

The aim of the study is to substantiate agent orchestration as the architectural and managerial foundation for intelligent BPM and to define the functions, controls, and performance indicators required for enterprise implementation.

Workflow automation and agent orchestration

Agent orchestration is not equivalent to chaining several model calls. It creates a managed process instance, decomposes the objective into work units, assigns specialized agents, maintains shared state, controls tool invocation, and resolves exceptions. The orchestrator remains responsible for sequence, deadlines, retries, escalation, and completion criteria, while agents remain responsible for bounded analytical or execution tasks. This separation prevents agent reasoning from replacing process control. The integration of multi-agent systems (MAS) with no-code architectures expands configurability for business users. A visual layer can define goals, transitions, rules, and approval thresholds, while agents dynamically select actions within these boundaries [2].

Research on large language model-based MAS identifies profiling, perception, self-action, interaction, and evolution as recurring system components, which must be translated into explicit enterprise roles and interfaces [3]. No-code tooling reduces the effort required to prototype agent workflows, but it does not remove the need for technical validation. Declarative agent specifications, reusable components, interactive debugging, and workflow evaluation improve development accessibility [4].

Enterprise automation still requires version control, test datasets, separation of environments, and change approval, since a visually configured process can produce the same operational risks as conventional software. The transition is most justified where a process combines stable objectives with variable inputs and requires several forms of expertise. Research on AI agent orchestration emphasizes the allocation of tasks among specialized agents and the integration of their outputs into a controlled process [5]. A comparable modular principle is visible in adaptive financial decision systems, where forecasting, optimization, transaction control, and governance are separated while remaining coordinated within one execution architecture [6].

Functional architecture of the orchestration layer

The orchestration layer should connect process logic with models, data, and enterprise applications without granting unrestricted autonomy. Each work unit requires an identifier, input schema, expected output, time limit, access profile, and completion status. The functional distribution proposed for an enterprise implementation is summarized in Table 1.

Table 1

Functional structure and controls of enterprise agent orchestration [2-8]

Layer	Primary function	Required control and indicator
Process intake	Create the process instance, classify context and define the expected result.	Input validation; complete context rate; rejected request rate.
Orchestrator	Decompose work, route tasks, maintain status, deadlines, retries and escalation.	State checkpoints; cycle time; retry and timeout rate.
Specialized agents	Perform analysis, retrieval, planning, validation or domain-specific execution.	Bounded role contract; task success rate; output quality.
Memory and knowledge	Store process state and retrieve approved corporate knowledge.	Source traceability; context accuracy; stale-data incidents.
Tools and integration	Read or update ERP, CRM, BPM, document and analytical systems.	Least privilege; transaction success; rollback rate.
Governance and human control	Apply policies, approve high-risk actions and preserve audit evidence.	Approval thresholds; intervention rate; policy violations.

Table 1 shows that agent quality is only one element of process performance. Reliable automation depends on the continuity of state, controlled handoffs, and validated writes to corporate systems. Monitoring at the level of an individual model cannot reveal duplicate transactions, cyclic calls, lost context, or incomplete process closure, which makes orchestration telemetry a separate management object. Big-data analysis in project management illustrates the same division between prediction and decision authority: classification, forecasting, clustering, and natural language processing can identify risks and resource deviations, while the process layer determines when a recommendation becomes an action [9]. In critical infrastructure, distributed agents also require layered integration from sensing and control to coordination and strategic decision support [8].

A financial application provides a compact example of coordinated model roles. The reviewed framework combines return and volatility forecasting, reinforcement learning-based allocation, transaction boundaries, and governance constraints. It reports an out-of-sample Sharpe ratio of 1.38, a 41% reduction in maximum drawdown during stress periods, and processing of a 50-asset portfolio in less than 25 ms [6]. These results are case-specific, yet they demonstrate that orchestration converts separate predictive modules into a measurable decision cycle rather than a set of disconnected model outputs.

Governance and efficiency assessment

Agentic BPM requires controls that begin before execution. Tool access should follow least-privilege rules; input and output schemas should be machine-verifiable; sensitive actions should require policy checks or human approval. Process logs must record the model, prompt, tool version, retrieved evidence, routing decision, and final transaction. The NIST generative AI profile supports lifecycle-based risk management, testing, documentation, and continuous monitoring rather than one-time model approval [7]. Operational evaluation should combine process and model metrics. A compact set includes cycle time, straight-through processing rate, exception frequency, human intervention, task completion, cost per completed case, rollback rate, security violations, and output quality.

A 2026 production study of end-to-end standard operating procedure automation reported 85-97% accuracy across seven categories, a reduction in case resolution time from 30 to 5 minutes, and deployment involving more than 1,000 account managers [10]. The results indicate that measurable value emerges when agent reasoning is supported by structured procedures, document verification, and reliable web or system interaction. Implementation should begin with processes whose outcomes are explicit and whose exceptions can be classified. A read-only pilot can validate retrieval, routing, and recommendations before write permissions are introduced. Scaling is justified after the system demonstrates stable quality under normal load, controlled degradation during model or integration failure, and reproducible human override. No-code configuration can accelerate adaptation, while production changes still require versioning, regression tests, and accountable ownership.

Conclusion

Agent orchestration forms the control foundation of intelligent BPM when enterprise work requires contextual interpretation, several specialized capabilities, and coordinated actions across digital systems. Its function is to preserve process identity, state, sequence, access boundaries, and auditability while agents perform bounded reasoning and execution tasks. The proposed architecture links process intake, orchestration, specialized agents, memory, integration, and governance. Operational efficiency should be evaluated through cycle time, completion quality, exception handling, cost, and control indicators rather than through agent autonomy alone. Applied evidence shows that coordinated models can support adaptive financial allocation, project analytics, critical-infrastructure decisions, and end-to-end operational procedures. Enterprise deployment remains appropriate only where autonomy is framed by explicit contracts, validated tools, failure recovery, and human responsibility for high-impact decisions. This configuration allows agent systems to extend BPM beyond rigid workflows without converting process flexibility into unmanaged operational risk

Список литературы:

1. Singla A., Sukharevsky A., Hall B., Yee L., Chui M. The state of AI in 2025: Agents, innovation, and transformation // McKinsey & Company. 2025. URL: https://www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai (accessed: 26.07.2026).

2. Hrebeniuk O. A model for integrating multi-agent systems and no-code architectures for intelligent automation of enterprise processes // Cold Science. 2026. No. 25. P. 25-34.

3. Li X., Wang S., Zeng S., Wu Y., Yang Y. A survey on LLM-based multi-agent systems: Workflow, infrastructure, and challenges // Vicinagearth. 2024. Vol. 1. Article 9. DOI: 10.1007/s44336-024-00009-2.

4. Dibia V., Chen J., Bansal G., Syed S., Fourney A., Zhu E., Wang C., Amershi S. AUTOGEN STUDIO: A no-code developer tool for building and debugging multi-agent systems // Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations. 2024. P. 72-79. DOI: 10.18653/v1/2024.emnlp-demo.8.

5. Hrebeniuk O. Business process automation using AI agent orchestration systems // Journal of Advanced Research in Technical Science. 2026. No. 52. P. 83-91.

6. Zharmagambetov Y.N. Dynamic asset allocation in unit-linked products: A machine learning approach to navigating market volatility // Professional Bulletin: Economics and Management. 2026. No. 1. P. 55-63.

7. Autio C., Schwartz R., Dunietz J., Jain S., Stanley M., Tabassi E., Hall P., Roberts K. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1. National Institute of Standards and Technology. 2024. DOI: 10.6028/NIST.AI.600-1.

8. Gvilava N.T. Autonomous intelligent agents in decision support systems for critical infrastructure // Professional Bulletin: Information Technology and Security. 2025. No. 1. P. 55-61.

9. Pshychenko D. Study of artificial intelligence models for big data analysis in project management // International Journal of Humanities and Natural Sciences. 2024. No. 8-3 (95). P. 180-185. DOI: 10.24412/2500-1000-2024-8-3-180-185.

10. Singh A., Agrawal S., Adhikari S., Puranik V.S., Tiwari S., Assudani D. Agent-Ops: A multi-agent orchestration framework for end-to-end SOP automation in e-commerce operations // Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics. Vol. 6: Industry Track. 2026. P. 436-446. DOI: 10.18653/v1/2026.acl-industry.29