

Artificial Intelligence Agents in ERP Ecosystems: An Analysis of the Gap Between Technological Discourse and Actual Integration Maturity

Octavian DOSPINESCU¹, Cristian SIMIONESCU², Daniel-Florin DĂNILOAIA³

¹Faculty of Economics and Business Administration, Alexandru Ioan Cuza University, Iași,
Romania

²Faculty of Computer Science, Alexandru Ioan Cuza University, Iași, Romania and Nexus
Media, Romania

³Faculty of Economics and Business Administration, Alexandru Ioan Cuza University, Iași
and Nexus Media, Romania

doctav@uaic.ro, cristian@nexusmedia.ro, daniel.daniloaia@nexusmedia.ro

Enterprise Resource Planning (ERP) systems serve as the operational backbone of modern companies, representing a field undergoing continuous technical and operational transformation and improvement. Although technological realities necessitate the modernization and integration of new technologies into ERP systems, the integration of autonomous artificial intelligence (AI) agents into these ecosystems currently lags significantly behind the marketing rhetoric surrounding the field. This article aims to highlight and analyze the discrepancy between the "agentic" rhetoric promoted by ERP vendors and the actual maturity level of autonomous functionalities implemented in market-leading platforms. Based on a systematic review of the literature from the past decade, we formulated three research hypotheses regarding: (H1) the difference between the prevalence of conversational AI assistants and that of AI agents with genuine transactional autonomy; (H2) the relationship between an ERP platform's cloud-native architecture and the maturity level of its integrated AI agents; and (H3) the gap between the intensity of "agentic" marketing rhetoric and the documented governance levels of these systems. The hypotheses were tested using a quantitative content analysis applied to the public documentation of 12 internationally representative ERP platforms. The analysis employed a coding instrument structured around four dimensions: the presence of conversational assistants within the ERP ecosystem, actual agentic maturity, the intensity of marketing rhetoric, and documented governance. The results of our objective analysis confirm the existence of a statistically significant gap ($p=0.002$) between truly autonomous functionalities and those that are merely assistive. At the same time, the analysis rejects ($p=0.433$) the existence of a significant association between cloud-native architecture and agentic maturity. Furthermore, our findings ($p<0.001$, mean gap = 1.44 points on a normalized 5-point scale) indicate a significant discrepancy between documented governance and marketing discourse. Our article makes a significant contribution to the study of agentic AI within the context of ERP systems, while also offering recommendations for managerial practice and potential avenues for future research.

Keywords: Agentic AI, ERP systems, Artificial intelligence agents, Digital transformation, Autonomous automation

DOI: 10.24818/issn14531305/30.3.2026.02

1 Introduction

Over the past three decades, ERP (Enterprise Resource Planning) systems have served as a fundamental infrastructure for medium- and large-sized organizations, primarily aiming to integrate operations and processes—spanning finance and accounting,

human resources, and production, as well as customer and supplier relationships. According to studies by [1] and [2], the core concept underpinning these systems was a unified data architecture. From an evolutionary perspective, this architecture was utilized in monolithic on-premises

implementations during the 1990s and 2000s—such as SAP R/3, Oracle E-Business Suite, and PeopleSoft—before advancing toward the cloud-native platforms of the last decade in response to persistent demands for scalability, efficiency, and integration [3], [4]. The most recent stage of this evolution involves the proliferation of generative artificial intelligence and autonomous AI agents within the extended architecture of ERP platforms, in continuation of early research into agent-based architectures for small and medium-sized enterprises [5] that began nearly 20 years ago. Against the backdrop of automation and the integration of AI capabilities, conceptual approaches to multi-agent systems [6] have given rise to philosophies centered on the concept of Extended Enterprise Resource Planning. According to [7], small and medium-sized enterprises (SMEs) face significant challenges in aligning IT infrastructure with business objectives while navigating digital transformation.

Specialized literature [8] highlights a clear distinction between assistive AI and agentic AI, a distinction essential for understanding the current state of the field. According to [9] and [10], assistive AI focuses on conversational copilots, content generation, and summarization, offering users suggestions and recommendations to inform their decision-making. Conversely, an intelligent agent is defined as an autonomous software entity capable of understanding context, planning, and executing a task or set of tasks to successfully achieve a specific goal. Within an integrated ERP system environment, an autonomous AI agent does not merely recommend an action (such as paying a supplier invoice by its due date or rebalancing inventory levels) but actually executes the transaction within the system. Furthermore, such autonomous execution must adhere to governance rules and, ideally, incorporate mechanisms for human oversight (human-in-the-loop). This approach aligns with the general direction of the multi-agent AI paradigm, which is increasingly being addressed in the scientific literature [11].

Since 2024, a marketing narrative centered on the term "agentic" has emerged among major ERP system providers. Recent research [4] indicates that Oracle has expanded Fusion Cloud ERP with AI Studio agents, while SAP has launched Joule Agents and Joule Assistants [12]. Similarly, Microsoft has enhanced Dynamics 365 by introducing an "agentic layer," and Workday has developed its own Illuminate Agents product, leveraging over a billion annual financial transactions. Analyses by specialized firms, based on data from official reports, confirm the accelerated pace of this transformation. According to Gartner's estimates [13], only 5% of ERP applications incorporated specific AI agents in 2025, but this figure is projected to rise to 40% by the end of 2026. Along these lines, McKinsey & Company [14] suggests that the effort required to implement enterprise applications could be reduced by up to 50% through the use of AI agents.

Research gap

Despite these developments—which are touted as having a major, far-reaching impact—the literature highlights a significant gap between the pace of market rhetoric and rigorous evidence regarding the actual level of functional autonomy possessed by AI agents within existing ERP systems. Indeed, very recent studies [15] on AI integration in ERP tend to focus on specific aspects—such as qualitative analyses of implementation challenges or individual research into a single functional module or traditional vendor [16], [12]. These disparate studies and research efforts fail to provide a systematic, quantitative, and comprehensive comparison of leading ERP solutions in terms of agentic maturity. At the same time, it is worth noting that most publicly available reports originate from consulting firms; while highly reputable, these reports often lack replicability and fail to employ a standardized methodology that would ensure full comparability across ERP platforms. This lack of methodological rigor creates a genuine risk of overestimating the maturity level of agentic AI within enterprise ecosystems. In the literature [17], this

phenomenon is termed "agent washing" and refers to the rebranding of existing components (such as Robotic Process Automation—RPA, chatbots, etc.) as "AI agents," despite the fact that these components lack substantial agentic capabilities.

Research goal

Given the research gap objectively identified through a review of the specialized literature, the objective of this study is to evaluate the actual level of representation and maturity of autonomous AI agents within major internationally available ERP ecosystems, in comparison to purely assistive functionalities. Furthermore, we aim to examine potential relationships between the architectural characteristics of ERP systems (specifically, cloud-native architecture) and actual agentic maturity. In doing so, we contribute to the body of knowledge by scientifically highlighting the discrepancy between marketing rhetoric and the governance of these systems. To achieve these research objectives, we will employ a quantitative methodology based on the content analysis of public documentation.

Research question

Based on the previously described research gap and research goal, we can define the research question that will guide the entire scientific inquiry of this article: "To what extent do contemporary ERP ecosystems genuinely integrate autonomous AI agents—as opposed to the rhetoric regarding agency promoted by vendors—and what structural factors (architecture, governance) explain the differences in maturity observed across platforms?" This research question allows us to break down the following specific questions, which we will test as research hypotheses in the subsequent sections of the study:

- a) Is there a significant difference between the prevalence of assistive AI capabilities and that of truly autonomous AI agents regarding actual trading?
- b) Is the modern cloud-native architecture of integrated ERP systems

associated with a higher level of agentic maturity?

- c) Is there a significant gap between the level of governance documented for enterprise systems and the intensity of "agentic" marketing discourse?

The article is structured into several sections: Section 2 outlines the rationale for the research hypotheses based on a detailed review of the literature, while Section 3 presents the research methodology; Section 4 details the results obtained from the quantitative analysis and hypothesis testing. Based on these findings, Section 5 discusses the theoretical and practical implications, and Section 6 summarizes the conclusions, research limitations, and potential future directions.

2 Literature Review and Hypotheses Development

2.1 The evolution of ERP systems from classic data integration to advanced intelligent orchestration

Chronologically, ERP systems emerged as applications designed to process data from across the entire organization in an integrated manner. In doing so, according to [1], they eliminated the shortcomings and inconveniences associated with approaches based on isolated data islands or silos. As data integration evolved, ERP systems also drove the standardization or normalization of business processes. From an evolutionary perspective, enterprise systems undergo several phases - documented in detail in the literature [3] as a specific lifecycle—progressing from the project phase to stabilization, and subsequently moving toward a stage of continuous improvement. Researchers specializing in ERP systems [2] have rightly observed that, driven by competitive pressures and increasingly sophisticated customer demands, these systems constantly evolve through the implementation of new functionalities, modules, and interconnection capabilities. Ultimately, this evolution aims to leverage raw data, transforming it into useful

information and knowledge that supports decision-making processes.

This evolution is highly palpable, particularly over the last 5–10 years, a period during which providers of "traditional" enterprise systems have developed solutions relying predominantly on cloud models and the Software-as-a-Service (SaaS) approach. These innovations have facilitated and accelerated the development and integration of modules incorporating artificial intelligence. Official reports [18] have projected the integrated systems market to reach a value of nearly \$65 billion by 2024, with a growth rate exceeding 11%. Forbes [19] identifies 2025 as an inflection point for the transition of AI agents into the stage of actual, practical use.

2.2 AI Agents – More Than Just Automation

The concept of an intelligent agent refers to a software system capable of understanding the context in which it operates and performing autonomous actions [9]. These actions aim to achieve a specific, clearly defined objective. AI agents also possess continuous perception and reasoning capabilities [10]. The emergence of Large Language Models (LLMs) has acted as a powerful catalyst for AI agents, as LLMs can serve as a "reasoning engine" for them. Furthermore, LLMs can break down complex problems into fundamental sub-problems [20], significantly enhancing the efficiency of AI agents. Moreover, intelligent agents are becoming a subject of interest within higher education as well [21].

In the specific context of enterprise applications, artificial intelligence tools exhibit varying levels of maturity: RPA, assistive, and agentic. RPA (Robotic Process Automation) tools typically execute highly structured, repetitive tasks but lack the adaptability to handle the exceptions that may arise [22]. Assistive tools support user decision-making by generating content, recommendations, and specific analysis results. The agentic maturity level entails a high degree of autonomy for the AI agent, as

it can make and effectively implement decisions independently. Agentic tools are superior to RPA tools because, when exceptions to the "normal" workflow occur, they can manage the situation using reasoning similar to that of a human [23]. Unlike the agentic approach, the RPA maturity level "stalls" when encountering unanticipated exceptions.

2.3 AI Agents in ERP Ecosystems

An analysis of the specialized literature and official documentation reveals that the ERP systems industry shows varying levels of interest in integrating AI agents into existing software. For instance, FinRobot is described by [16] as a system where financial and accounting processes are automated through an AI agent architecture. Another example is the Joule agents within the SAP system, which collaborate in areas such as logistics, procurement, and finance. However, despite being advanced tools, some authors [12] indicate that these SAP agents remain quite limited in terms of operational capabilities.

The situation varies across specific industries. For instance, the research conducted by [24] indicates that in the manufacturing sector, most AI-driven systems function merely as assistants rather than advancing to the level of autonomous execution. In the healthcare sector, a study shows that process governance poses a major challenge regarding ERP-level integration. The challenges of integrating AI into enterprise systems are manifold and are highlighted in the literature [25], [15]: data quality, high costs associated with integration with legacy systems, and a lack of interoperability standards between agents. Employee resistance to change is also a factor that cannot be overlooked.

In terms of satisfaction and productivity, the results obtained by [26] show increases of 30% and 25%, respectively, when predictive analytics provided by AI systems are utilized; however, these figures do not include autonomous agentic execution.

Given these aspects, we propose the following research hypothesis:

H1: The proportion of major ERP platforms featuring purely assistive AI capabilities is significantly higher than the proportion of platforms featuring AI agents with genuine transactional autonomy.

Analyses in the specialized literature [18], [4] tend to suggest that integrating new AI agents with existing "classic" ecosystems is more easily achieved within cloud-native architectures. The key elements consist of microservices and open APIs. These arguments and trends lead us to formulate the following research hypothesis:

H2: There is a significant positive association between the cloud-native (SaaS) architecture of an ERP platform and the maturity level of its integrated AI agents.

At the same time, however, consulting firms [14] advise that the transformations associated with AI agents should not be fragmented but rather implemented at a functional level. A recent survey [13] highlights that while only 17% of organizations have actually implemented AI agents, over 60% have expressed interest in adopting such technology within the next two years. This combination of conflicting figures—showing optimistic medium-term projections or trends contrasted with very low actual implementation rates—prompts us to formulate research hypotheses for scientific analysis. Furthermore, [17] points to the prevalence of "agent washing" - the rebranding of existing AI assistants or RPA tools as "agents" despite lacking genuine agentic capabilities. This suggests a marketing initiative detached from actual operational reality.

For this reason, we aim to test the following research hypothesis:

H3: There is a significant gap between the intensity of the "agentic" marketing rhetoric of ERP vendors and the publicly documented governance level for the respective AI agents (human oversight mechanisms, auditing, explainability).

Through these three research hypotheses, we aim to develop a robust research model that takes into account function 1 maturity, the architectural factors of enterprise systems, and the communication gap.

3 Research Methodology

3.1 Research Design

To test the proposed hypotheses and research model, we employed a quantitative design. A distinctive feature of our design is that the quantitative analysis utilizes secondary data from scientific articles, official reports, and public technical documentation. Consequently, our methodology is not affected by the subjectivity inherent in questionnaire-based analyses. The methodology we adopted is recognized and well-established in the specialized scientific literature [27] regarding enterprise systems. It enables the systematic comparison of a large number of cases, offering the objective possibility of replication. Furthermore, this methodology has the advantage of not depending on factors such as sampling and respondent availability.

3.2 Research Sample

To conduct an objective analysis of the research hypotheses, we selected 12 integrated ERP platforms that are globally representative. This selection was based on a review of the specialized literature as well as specific market shares [28]. Consequently, the complete list of enterprise systems included in this study comprises: SAP (with Joule agents), Oracle Fusion Cloud ERP, Microsoft Dynamics 365, Workday Illuminate, InforCloudSuite, Oracle NetSuite, Odoo, Sage X3/Intacct, IFS Cloud, Epicor Kinetic, Acumatica Cloud ERP, and Unit4 ERPx. The selection process also relied on professional judgment to ensure the list included both "Tier 1" products (Microsoft, SAP, Oracle) and vendors targeting the small and medium-sized enterprise segment (Acumatica, Odoo, Unit4). In this way, we ensured sufficient representativeness to test research hypothesis H2 regarding architecture.

3.3 Coding Instrument

Given the authors' experience in the field of enterprise systems and artificial intelligence applications [29], [30], [31], [32], we employed a tool structured around four dimensions, utilizing documentation that was publicly available at the time this article was written:

- **Marketing_Score:** a dimension quantified on an ordinal scale (0–5). This dimension measures the intensity of "agentic" marketing rhetoric in the software vendor's official materials. Coding was based on the frequency and prominence of terms such as "agent," "agentic," and "autonomous" in key official public presentations and materials.
- **Assist_AI:** a binary dimension (0/1). This dimension quantifies the presence of a conversational assistant based on generative artificial intelligence. This quantification is performed regardless of the assistant's level of autonomy or the accuracy of the responses it provides.
- **Governance_Score:** a dimension based on an ordinal scale (0–3). This dimension highlights the level of publicly documented governance. The associated values are: 0 – undocumented; 1 – basic logging/auditing; 2 – human approval points that are explicitly documented and proceduralized; 3 – full auditing, documented explainability mechanisms, and specific roles and role-based controls.
- **Agent_Maturity:** a dimension based on an ordinal scale (0–5). This dimension measures the level of agent maturity, with the following value definitions: 0 – no AI functionality; 1 – chatbot or assisted search, but no action taken; 2 – AI-assisted generation of reports or other content, but no autonomous execution; 3 – automation of AI-recommended tasks, but requiring prior human approval; 4 – autonomous agent with "human-in-the-loop" control

points within a single functional module; 5 – multi-agent orchestration across multiple functional modules, with minimal human intervention.

In addition to these dimensions, we also coded the **Cloud_Native** variable with the following values: 0 – an initial on-premises architecture subsequently migrated to the cloud; 1 – a SaaS architecture from the moment the system was launched into production.

3.4 Methods of Statistical Analysis

Given the specific nature of the research hypotheses and the coding schemes used for the variables under analysis, distinct methods were employed for each case.

For hypothesis H1, the McNemar test was used, as it is the most appropriate method for comparing two correlated binary variables.

To test hypothesis H2, the relatively small sample size ($n = 12$ products) and the ordinal nature of the **Agent_Maturity** variable were taken into account. Consequently, both the Pearson correlation coefficient and the Kendall test were applied to compare the distribution of agentic maturity between cloud-native and non-cloud-native platforms. For hypothesis H3, the governance and marketing scores were normalized to bring both onto a common 0–5 scale; specifically, the governance score was rescaled from the original 0–3 range to the 0–5 range. Following normalization, the Wilcoxon signed-rank test was applied to the marketing and governance scores.

All tests were conducted and interpreted using the conventional significance level of $\alpha = 0.05$.

4 Results

The secondary data underpinning our analysis are presented in the following table. As previously detailed, these scores cover the top 12 enterprise platforms globally.

Table 1. Supporting data for quantitative analysis

ERP platform	Assist_AI	Agent_Maturity	Governance_Score	Marketing_Score	Cloud_Native	Launch_year
SAP S/4HANA (Joule)	1	4	2	5	0	2024
Oracle Fusion Cloud ERP	1	4	2	5	0	2024
Microsoft Dynamics 365	1	3	2	4	0	2024
Workday (Illuminate)	1	3	1	4	1	2024
Infor CloudSuite	1	2	1	3	0	2024
Oracle NetSuite	1	2	1	3	1	2023
Odoo	1	1	0	2	1	2025
Sage Intacct/X3	1	2	1	3	0	2024
IFS Cloud	1	2	1	3	0	2024
Epicor Kinetic	1	2	1	3	0	2024
Acumatica	1	2	1	2	1	2024
Unit4 ERPx	1	1	0	2	0	2025

An analysis of the data in Table 1 reveals that all 12 enterprise platforms feature a conversational assistant based on generative artificial intelligence. While this study does not evaluate the quality or performance of the conversational assistants for each individual application, it is evident that assistive functionality has become a „de facto” standard as of 2026.

On the other hand, the scores in Table 1 indicate that the average level of agentic maturity (Agent_Maturity) is relatively low—standing at just 2.33 points (standard deviation = 0.98) on a scale of 0 to 5. It is worth noting that, at present, none of the analyzed tools achieve the maximum score of 5, which corresponds to fully autonomous, cross-modular multi-agent orchestration. Only Oracle Fusion Cloud ERP and SAP S/4HANA with Joule reach level 4 (representing an autonomous agent with documented human checkpoints).

An analysis and interpretation of the average documented governance score (Governance_Score) reveals a mean value of 1.08 points with a standard deviation of 0.67. Proportionally, on a scale of 0 to 3, this score is significantly lower than the average marketing score (Marketing_Score), which stands at 3.25 points (standard deviation of 1.06) on a scale of 0 to 5. This discrepancy serves as the initial descriptive indicator of the gap that will be tested under hypothesis H3.

4.1 Testing hypothesis H1

Hypothesis H1, stated earlier, posits that the proportion of major ERP platforms featuring purely assistive AI capabilities is significantly higher than the proportion of platforms featuring AI agents with genuine transactional autonomy. In our representative sample, there are $n=12$ platforms with assistive capabilities and 2 platforms reaching a level of agentic maturity (Agent_Maturity ≥ 4). We applied the McNemar test to the discordant pairs ($b = 10$ platforms with an AI assistant but no genuine autonomous agent, and $c = 0$ platforms with a genuine autonomous agent but no AI assistant), and the resulting statistic is significant ($p = 0.03$).

Based on this result, we can state that **hypothesis H1 is confirmed**, meaning there is a statistically significant difference between the prevalence of assistive AI functionalities and that of AI agents with genuine transactional autonomy. Assistive AI functionalities are nearly universal in the market, whereas the prevalence of autonomous agents remains marginal (less than 17%).

The quantitative result we obtained provides direct, objective confirmation of the qualitative observations made by [19]. Furthermore, our result aligns with projections made by [13], which estimated an adoption rate of less than 5% for task-specific agents by 2025.

4.2 Testing hypothesis H2

Under hypothesis H2, we tested the association between cloud architecture (Cloud_Native) and agentic maturity (Agent_Maturity). To test H2, we calculated the Pearson correlation coefficient and obtained a value of $r = -0.25$ ($p = 0.433$). Since the p -value exceeds the 0.05 threshold, we conclude that hypothesis H2 is not confirmed. We also applied the Kendall test to the data in Table 1; the results likewise indicated the absence of a significant association between the cloud-native platform group and the non-cloud-native platform group. The test values are: Kendall τ -b = -0.204 , $p = 0.466$. The absence of a significant association is supported by the p -value, which is higher than the conventional 0.05 threshold. Consequently, **hypothesis H2 is rejected** based on the econometric data obtained from the analysis.

An examination of the descriptive data reveals that the cloud-native platforms in the sample actually exhibit a slightly lower—rather than higher—average score regarding the agentic maturity indicator.

In this context, we also conducted an additional exploratory analysis of the relationship between Agent_Maturity and Governance_Score, obtaining a strong, significant, and positive Pearson correlation: $r = 0.92$, $p < 0.001$. These values suggest that platforms with a higher degree of agentic maturity tend to document more rigorous governance mechanisms. This result is consistent with best practice recommendations found in the literature [15].

4.3 Testing hypothesis H3

Hypothesis H3 aims to test the gap between the marketing score (*Marketing_Score*) and the documented governance score (*Governance_Score*). We reiterate that the *Marketing_Score* uses a scale of 0 to 5, and the *Governance_Score* was renormalized to the same scale.

The Wilcoxon signed-rank test yielded a highly statistically significant result (statistic = 0.00, $p = 0.0005$). All 12 platforms in the sample exhibit a higher marketing score than the normalized governance score. The mean

gap between the two scores is 1.44 points (SD = 0.56) on the 5-point scale, equivalent to a relative gap of approximately 44% of the scale's maximum range.

Given these findings, it is evident that **hypothesis H3 is confirmed**. Thus, we can state with certainty that there is a significant gap between the intensity of the "agentic" marketing rhetoric used by enterprise system providers and the publicly documented governance level of the respective AI agents. The quantitative result obtained through our econometric analysis provides an objective basis for the warning issued by [17] regarding the phenomenon of "agent washing."

5 Discussions

The results obtained from testing the research model reveal a nuanced reality regarding the degree of agentic integration within Enterprise Resource Planning (ERP) systems. The current state of AI agent integration in ERP systems lies somewhere at the intersection of marketing hype and verifiable technical reality. By confirming hypothesis H1, we provide an econometric quantification of an observation that has not previously been systematically measured in the literature. Thus, although generative assistive AI technology has become an industry standard for enterprise systems, the capability for genuine transactional autonomy remains the exclusive domain of a select group of major vendors (such as SAP and Oracle)—companies possessing substantial resources for research and development. Our findings align with the projections made by [13] and the qualitative results presented by [24] regarding the prevalence of decision support over autonomous execution in ERP systems. Our results also align with those obtained by [33] regarding the state of multi-agent collaboration in integrated ERP environments. Thus, currently, most agents possess LLM-like capabilities, but they typically do not leverage the benefits of multi-agent collaboration, such as divergent thinking or reasoning.

We consider the rejection of hypothesis H2 to be a highly significant result of our research,

as it refutes an intuitive assumption found in the literature—namely, that cloud-native architecture acts, in itself, as a catalyst for agentic maturity. Somewhat paradoxically, the data suggest that the determining factor for agentic maturity appears to be the software vendor's market dominance—and potentially the scale of its R&D investment—rather than the underlying technical architecture. For instance, SAP and Oracle, whose systems rely on on-premises architectures subsequently migrated to the cloud, exhibit the highest agentic maturity scores. This aspect is also documented in the specialized literature [34] and [35], which highlights the importance of specialized frameworks and specific adaptations. Conversely, smaller vendors offering platforms that were cloud-native from the outset (such as Acumatica and Odoo) remain at lower levels of maturity. However, we believe this result should be interpreted with caution, given the relatively limited size of the sample analyzed (12 software products). Nevertheless, this initial finding paves the way for future research into how factors such as organizational resources and a vendor's position in the enterprise systems market might influence agentic maturity.

Confirming hypothesis H3 highlights significant managerial implications. Consequently, companies wishing to implement enterprise systems with agentic functionalities must pay close attention to the discrepancies between marketing communications and the actual, documented level of governance. Interested companies are advised to first gather data based on precise figures regarding aspects such as access control, auditing, explainability, and points of human intervention.

Viewed through the lens of literature [3] analyzing the adoption lifecycle of integrated systems, our study results suggest that, as of 2026, the integration of AI agents into enterprise systems is in a relatively early, "project-based" phase. This stage is characterized by marketing efforts and messaging that overstate actual agentic functionalities and capabilities. Furthermore, there is currently no standard for AI agent

integration; the initiatives undertaken by integrated system vendors amount to little more than a collection of disparate experiments.

6 Conclusions

6.1 Main findings

This study employed a quantitative methodology based on content analysis to investigate the actual level of integration of autonomous AI agents within enterprise ecosystems, contrasting this with the associated marketing discourse. The content analysis covered 12 ERP platforms representative of the integrated systems market. In summary, the findings relate to the three research hypotheses tested. The results revealed that purely assistive AI functionalities have become a „de facto” market standard, appearing in 100% of the analyzed platforms. Conversely, AI agents possessing genuine transactional autonomy remain marginal, as they appear at a high level of maturity in only 16.7% of the cases analyzed. This disparity was statistically confirmed through the testing of hypothesis H1. Another finding is that the cloud-native architecture of an ERP platform is not associated with a higher level of agentic maturity. Consequently, hypothesis H2 was rejected, contradicting an intuitive assumption derived from the existing literature.

A key finding highlighted by the testing of hypothesis H3 reveals a significant difference between the intensity of "agentic" marketing discourse and the publicly documented level of governance for AI agents within the analyzed enterprise environments.

Overall, our results confirm the existence of a significant and quantifiable gap between the agentic rhetoric surrounding enterprise systems and the documentable technical level of current implementations.

6.2 Limitations

Objectively, our research presents certain limitations. The first concerns the small size of the sample, which consists of 12 integrated ERP system products and their respective vendors. While we have argued that these 12

platforms are representative of the global market in terms of market share and economic impact, the statistical power of the tests we applied is limited. Consequently, our results must be interpreted with this limitation in mind. We believe the small sample size has a particularly limiting effect on hypothesis H2 within our research model.

Secondly, the content analysis methodology—while highly rigorous and fully replicable—is constrained by the timeframe in which the actual analysis was conducted (specifically, May–June 2026). We highlight this because content analysis relies on the accuracy and completeness of data provided by official enterprise system vendors. Moreover, the field is highly dynamic and rapidly evolving, meaning the content in question is subject to frequent changes.

Thirdly, the scores regarding agentic maturity and marketing intensity entail a level of interpretive judgment on the part of the evaluators, even though they are experienced in the field of ERP systems.

Fourthly, our research does not cover aspects related to the agents' actual performance in concrete operational situations. As outlined in the article, empirical evaluation metrics—such as the success rate of autonomously executed transactions, the frequency of corrective human interventions, costs associated with agent execution, etc.—were not analyzed.

6.3 Future research directions

Given the results obtained and the previously identified objective limitations, we propose several potential future research directions. Clearly, a primary direction involves expanding the sample to include other enterprise software products. This would enhance the representativeness of the findings, allowing for a more confident extrapolation of the results.

A second research avenue involves periodically repeating the content analysis to conduct a longitudinal study, thereby highlighting the evolution over time of the actual maturation rate of agentic AI within

ERP ecosystems. This analysis could also incorporate a comparison with projections from market research firms.

Thirdly, the study presented in this article could be extended by supplementing the content analysis with actual operational data. This would allow for a deeper empirical examination of the relationship between actual and documented/stated agentic maturity by considering metrics such as error rates, process speed, cost reductions, and savings in time and human resources. We anticipate that this proposed research direction would require partnerships with companies that have already implemented AI agents within their ERP systems, enabling access to real-world data from agentic work environments.

A fourth research direction could be an objective analysis of the interoperability of agents within integrated ERP systems sourced from different vendors. Interoperability mechanisms are closely linked to the concept of standardization regarding AI agents in enterprise ecosystems.

Building on the objective findings of this study—which highlight a clear trend toward agent-based environments and capabilities within integrated ERP systems—a fifth, complementary research direction could involve analyzing the perceptions and readiness of the workforce regarding working and collaborating with AI agents. This represents an organizational dimension that extends the study presented in this article.

References

- [1] T. Davenport, "Putting the Enterprise into the Enterprise System," *Harvard Business Review*, vol. 76, no. 4, pp. 121-131, 1998.
- [2] Y. Moon, "Enterprise Resource Planning (ERP): a review of the literature," *International Journal of Management and Enterprise Development*, vol. 4, no. 3, pp. 235-264, 2007.
- [3] M. Markus and C. Tanis, "The Enterprise System Experience—from Adoption to Success," in *Framing the Domains of IT Management: Projecting the Future through the Past*, Cincinnati, Pinnaflex

- Educational Resources Inc, 2000, pp. 173-207.
- [4] S. Padur, "The Future of Enterprise ERP Modernization with AI: From Monolithic Systems to Generative, Composable, and Autonomous Platforms," *Journal of Artificial Intelligence, Machine Learning and Data Science*, vol. 3, no. 1, pp. 1-3, 2025.
- [5] R. Fu, Z. Xin and C. Chen, "Research on Agent-Based Architecture of ERP for Small & Medium-Size Enterprise," in *2007 International Conference on Machine Learning and Cybernetics*, Hong Kong, China, 2007.
- [6] D. Plikynas, "Extended enterprise resource planning: conceptual approach using multiagent systems," in *ACMOS'08: Proceedings of the 10th WSEAS International Conference on Automatic Control, Modelling & Simulation*, Istanbul, Turkey, 2008.
- [7] A. Pingilili, N. Letsie, G. Nzimande, B. Thango and L. Matshaka, "Guiding IT Growth and Sustaining Performance in SMEs Through Enterprise Architecture and Information Management: A Systematic Review," *Businesses*, vol. 5, no. 2, p. 17, 2025.
- [8] Z. Liu, T. Wang, X. Wang and L. Wang, "Agentic ERP: Multi-Agent Large Language Model Architecture for Autonomous Enterprise Resource Planning," *arXiv preprint*, 2026.
- [9] M. Wooldridge and N. Jennings, "Intelligent agents: theory and practice," *The Knowledge Engineering Review*, vol. 10, no. 2, pp. 115-152, 1995.
- [10] S. Russell and P. Norvig, *Artificial Intelligence - A Modern Approach*, U.K.: Pearson Education Limited, 2021.
- [11] S. Allmendinger, L. Bonenberger and K. Endres, "Multi-agent AI," *Electronic Markets*, vol. 36, no. 18, 2026.
- [12] D. Eyo, O. Adeniran and K. Osobase, "Agentic AI in SAP: Collaborative joule agents across procurement, finance and logistics," *Global Journal of Engineering and Technology Advances*, vol. 24, no. 2, pp. 91-108, 2025.
- [13] Gartner, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026, Up from Less Than 5% in 2025," 26 August 2025. [Online]. Available: <https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025>. [Accessed 14 June 2026].
- [14] McKinsey, "Bridging the great AI agent and ERP divide to unlock value at scale," 9 January 2026. [Online]. Available: <https://www.mckinsey.com/capabilities/mckinsey-technology/our-insights/bridging-the-great-ai-agent-and-erp-divide-to-unlock-value-at-scale>. [Accessed 15 June 2026].
- [15] M. Benjelloun, H. Hmamed, B. Rzine and A. Dadda, "Navigating Challenges When Integrating Artificial Intelligence with Enterprise Resource Planning: A Literature Review," in *International Conference on Advanced Intelligent Systems for Sustainable Development (AI2SD 2024)*. AI2SD 2024, Cham, 2025.
- [16] H. Yang, L. Lin, Y. She, X. Liao, J. Wang, R. Zhang, Y. Mo and C. Wang, "FinRobot: Generative Business Process AI Agents for Enterprise Resource Planning in Finance," *arXiv*, 2025.
- [17] Gartner, "Gartner Predicts Over 40% of Agentic AI Projects Will Be Canceled by End of 2027," 25 June 2025. [Online]. Available: <https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>. [Accessed 12 May 2026].
- [18] Saventech, "AI-Native ERP Systems in 2025: The Future of Enterprise Software," 2025. [Online]. Available: <https://saventech.com/ai-native-erp-systems-in-2025-the-future-of-enterprise-software/>. [Accessed 20 June 2025].
- [19] R. Kramer, "ERP Leveled Up In 2024 — Where Is It Headed In 2025?," *Forbes*,


- vol. 11 feb 2025, 2025.
- [20] M. Olariu, V. Buinceanu, C. Simionescu, O. Dospinescu, R. Georgescu, C. Tudor, A. Iftene and A. Bores, "RO-FIN-LLM: A Benchmark with LLM-as-a-Judge and Human Evaluators for Romanian Tax and Accounting," *Systems*, vol. 14, no. 3, p. 244, 2026.
- [21] O. Dospinescu, S. Buraga and N. Dospinescu, "Students' Perceptions of the Use of Artificial Intelligence Tools in Educational Activities," *Systems*, vol. 14, no. 6, p. 633, 2026.
- [22] A. Maldareanu and V. Dumitru, "The Integration of ERP Systems and Technologies for Process Automation: An Analysis of the Literature," *Transformations in Business & Economics*, vol. 24, no. 3, pp. 286-303, 2025.
- [23] Johari, J., "AI Agent Adoption in 2026: What the Data Shows," 20 February 2026. [Online]. Available: <https://joget.com/ai-agent-adoption-in-2026-what-the-analysts-data-shows/>. [Accessed 17 July 2026].
- [24] E. Niederwieser, D. Siegele and D. Matt, "AI-Driven ERP Systems: Integrating Large Language Models for Enhanced Customer Interaction and Operational Efficiency," *Zeitschrift für wirtschaftlichen Fabrikbetrieb*, vol. 120, no. s1, pp. 112-117, 2025.
- [25] M. Ghane, S. Sorooshian and M. Ang, "Artificial Intelligence agents and autonomous decision-making in business: a review," *Cogent Business & Management*, vol. 13, no. 1, 2026.
- [26] S. Mhaskey, "Integration of Artificial Intelligence (AI) in Enterprise Resource Planning (ERP) Systems: Opportunities, Challenges, and Implications," *International Journal of Computer Engineering in Research Trends*, vol. 11, no. 12, pp. 1-9, 2024.
- [27] K. Krippendorff, *Content analysis*, SAGE Publications Inc., 2019.
- [28] R. Jaimes, "Global ERP Statistics 2026: Market Size And Adoption Trends," 27 May 2026. [Online]. Available: <https://www.quantumrun.com/consulting/global-erp-statistics/>. [Accessed 11 July 2026].
- [29] O. Dospinescu, S. Buraga and N. Dospinescu, "Students' Perceptions of the Use of Artificial Intelligence Tools in Educational Activities," *Systems*, vol. 14, no. 6, p. 633, 2026.
- [30] M. Olariu, V. Buinceanu, C. Simionescu, O. Dospinescu, R. Georgescu, C. Tudor, A. Iftene and A. Bores, "RO-FIN-LLM: A Benchmark with LLM-as-a-Judge and Human Evaluators for Romanian Tax and Accounting," *Systems*, vol. 14, no. 3, p. 244, 2026.
- [31] A. Măldăreanu, V. Dumitru and O. Dospinescu, "An Analysis of User Reviews of ERP Systems in Small and Medium-Sized Enterprises," *Amfiteatru Economic*, vol. 28, no. 72, pp. 470-488, 2026.
- [32] C. Simionescu, A. Siriteanu, D. Dăniloia, O. Dospinescu, M. Olariu and A. Iftene, "ERP users intention to use artificial intelligence and automation instruments in Romanian businesses," *Cogent Business & Management*, vol. 13, no. 1, 2026.
- [33] J. Schnepf, M. Schwarz, B. Scheuermann and S. Anderer, "A Study on Multi-agent Collaboration for Business Process Automation in Enterprise Resource Planning Systems," in *Computational Intelligence. IJCCI 2025. Communications in Computer and Information Science*, 2026.
- [34] S. Sarferaz, "Implementing Agentic AI Into ERP Software," *IEEE Access*, vol. 13, pp. 178945-178960, 2025.
- [35] G. Stefanov and V. Kisimov, "Multi-Agent System for Dynamic Business KPI Selection, Evaluation and Quantification Based on Oracle EBS," *Future Internet*, vol. 18, no. 5, p. 268, 2026.



Octavian DOSPINESCU graduated the Faculty of Economics and Business Administration in 2000 and the Faculty of Informatics in 2001. He achieved the PhD in 2009 and he has published as author or co-author many articles. He is author and co-author of 10 books and teaches as full professor in the Department of Accounting, Information Systems and Statistics of the Faculty of Economics and Business Administration, University Alexandru Ioan Cuza, Iasi. He is interested in mobile devices software, enterprise resources planning systems, computer programming and decision support systems.



Cristian SIMIONESCU has a Ph.D. in Computer Science and he is Chief Technology Officer at Nexus Media. Being specialized in machine learning and artificial intelligence, he taught classes of data structures, algorithm design, probabilistic programming and modeling and software engineering at Faculty of Computer Science, University Alexandru Ioan Cuza, Iasi. He is a member of the Pionier AI implementation team, an artificial intelligence extension of the Nexus ERP software.



Daniel-Florin DANILOAIA is Ph.D. in Management and he is Associate Lecturer at the Faculty of Economics and Business Administration, Alexandru Ioan Cuza University of Iași. He has expertise as a research expert in national and European projects, focusing on management, artificial intelligence, and organizational development. He is also a researcher at Nexus Media in Artificial Intelligence.